



**DIRECTORY SERVICES**  
**SECURITY TECHNICAL IMPLEMENTATION GUIDE**  
Version 1, Release 1

24 August 2007

**Developed by DISA for the DoD**

UNCLASSIFIED

## **Trademark Information**

Active Directory, Microsoft, Windows, Windows NT, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Linux is a registered trademark of Linus Torvalds.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries.

UNIX® is a registered trademark of The Open Group.

All other names are registered trademarks or trademarks of their respective companies.

## TABLE OF CONTENTS

	Page
<b>1. INTRODUCTION.....</b>	<b>1</b>
1.1 Background.....	2
1.2 Authority.....	3
1.3 Scope.....	4
1.4 Writing Conventions.....	4
1.5 Vulnerability Severity Code Definitions.....	4
1.6 DISA Information Assurance Vulnerability Management (IAVM).....	6
1.7 STIG Distribution.....	6
1.8 Document Revisions.....	6
<b>2. DIRECTORY SERVICE ELEMENTS.....</b>	<b>7</b>
2.1. Introduction.....	7
2.2. Common Directory Service Elements.....	7
2.2.1. Directory Server.....	8
2.2.1.1. Authentication.....	8
2.2.1.2. Account Definitions.....	9
2.2.1.3. Query and Update Access.....	10
2.2.1.4. Other Server Issues.....	11
2.2.2. Directory Database.....	12
2.2.3. Replication.....	14
2.2.4. Administration Tools and Accounts.....	15
2.2.5. Network Ports and Protocols.....	16
2.3. Directory Synchronization Tools and Technology.....	17
<b>3. DIRECTORY SERVICE SECURITY REQUIREMENTS.....</b>	<b>21</b>
3.1. Security Design and Configuration.....	22
3.1.1. Product Design.....	22
3.1.2. Configuration and Implementation Integrity.....	23
3.1.3. Network Services.....	25
3.1.4. Software Integrity.....	26
3.1.5. Security Service Partitioning.....	29
3.2. Identification and Authentication.....	30
3.2.1. Password-Based Authentication.....	31
3.2.2. PKI-Based Authentication.....	33
3.3. Enclave and Computing Environment.....	33
3.3.1. Specific Content.....	34
3.3.2. Architecture and Cross-Directory Authentication.....	34
3.3.3. Data Access Control - Files.....	35
3.3.4. Data Access Control - Directory Database Objects.....	37
3.3.5. Data Change Auditing.....	39
3.3.6. Group Membership and Limiting Privileges.....	41
3.3.7. Functional Configuration.....	44
3.3.8. Data Transmission Confidentiality and Integrity.....	47
3.4. Enclave Boundary Defense.....	53
3.5. Physical and Environmental.....	55

3.6. Continuity .....	55
3.7. Vulnerability and Incident Management .....	58
<b>APPENDIX A. RELATED PUBLICATIONS .....</b>	<b>61</b>
<b>APPENDIX B. LIST OF ACRONYMS .....</b>	<b>65</b>
<b>APPENDIX C. ACTIVE DIRECTORY SPECIFIC ELEMENTS AND REQUIREMENTS</b>	
<b>69</b>	
C.1. Introduction.....	69
C.2. Active Directory Security Background.....	70
C.2.1. Active Directory Functional Level Considerations .....	71
C.2.2. Forest and Domain Architecture.....	72
C.2.2.1. Domains, Trees, and Forests .....	72
C.2.2.2. Replication, Sites, GC Servers, and FSMO Servers .....	74
C.2.2.3. Service Dependencies and AD Data Files .....	77
C.2.3. Group Memberships .....	80
C.2.3.1. SID Assignment and Use .....	80
C.2.3.2. Special Privileged Groups.....	81
C.2.3.3. Universal Groups .....	82
C.2.3.4. Group Nesting and Permission Strategies.....	82
C.2.3.5. OU Design .....	83
C.2.3.6. AD Object Quotas.....	83
C.2.4. Trust Relationships .....	84
C.2.4.1. Trust Properties and Terms .....	84
C.2.4.2. Automatically Defined Trusts.....	86
C.2.4.3. Manually Defined Trusts .....	86
C.2.5. Group Policy .....	89
C.2.5.1. Group Policy Components.....	89
C.2.5.2. Default GPOs .....	90
C.2.5.3. Application of GPOs to Objects .....	90
C.2.5.4. Group Policy Management Console (GPMC) and Group Policy Results Tool....	91
C.2.5.5. GPO Auditing and Backup .....	92
C.2.6. Ports and Protocols .....	92
C.3. Technology-Specific Security Requirements .....	95
C.3.1 Security Design and Configuration .....	95
C.3.2 Identification and Authentication .....	97
C.3.3 Enclave and Computing Environment.....	98
C.3.3.1. Architecture and Cross-Directory Authentication (AD Trusts).....	98
C.3.3.2. Data Access Control - AD Files.....	103
C.3.3.3. Data Access Control - AD Database Objects .....	104
C.3.3.4. Data Change Auditing - AD .....	106
C.3.3.5. Group Membership and Limiting Privileges - AD .....	109
C.3.3.6. Functional Configuration - AD.....	112
C.3.4 Physical and Environmental .....	113
C.3.5 Continuity .....	113
<b>APPENDIX D. ACTIVE DIRECTORY APPLICATION MODE SPECIFIC ELEMENTS</b>	
<b>AND REQUIREMENTS.....</b>	<b>115</b>
D.1. Introduction.....	115

D.2. Active Directory Application Mode Security Background .....	115
D.3. Technology-Specific Security Requirements .....	116
D.3.1. Enclave and Computing Environment.....	116
D.3.1.1. Data Access Control - ADAM Files .....	116
D.3.1.2. Data Change Auditing - ADAM.....	117
D.3.1.3. Group Membership and Limiting Privileges - ADAM.....	117
D.3.2. Continuity .....	117
<b>APPENDIX E. RED HAT DIRECTORY SERVER SPECIFIC ELEMENTS AND</b>	
<b>REQUIREMENTS.....</b>	<b>119</b>
E.1. Introduction.....	119
E.2. Red Hat Directory Server Security Background.....	119
E.2.1. Account and Group Considerations.....	120
E.2.1.1. Accounts .....	120
E.2.1.2. Groups.....	121
E.2.1.3. Account and Group-Related Controls.....	122
E.2.2. Object Access Control .....	124
E.2.3. Other Implementation Features and Details .....	127
E.2.3.1. Replication .....	127
E.2.3.2. Administrative Interfaces.....	128
E.2.3.3. Plug-ins .....	129
E.2.3.4. Gateways.....	129
E.2.3.5. Chaining.....	130
E.2.3.6. Ports and Protocols .....	130
E.3. Technology-Specific Security Requirements .....	131
E.3.1. Security Design and Configuration .....	132
E.3.2. Identification and Authentication .....	134
E.3.3. Enclave and Computing Environment.....	136
E.3.3.1. Cross-Directory Authentication (Pass-through Authentication).....	136
E.3.3.2. Data Access Control - RHDS Files.....	137
E.3.3.3. Data Access Control - RHDS Directory Database Objects .....	137
E.3.3.4. Data Change Auditing - RHDS.....	142
E.3.3.5. Group Membership and Limiting Privileges - RHDS .....	144
E.3.3.6. Functional Configuration - RHDS .....	145
E.3.3.7. Data Transmission Confidentiality and Integrity - RHDS.....	147
E.3.4. Continuity .....	150

## LIST OF FIGURES

Figure C-1. Sample AD Forest .....	73
Figure C-2. Sample Trusts .....	85
Figure E-1. RHDS Tree .....	124

## LIST OF TABLES

Table C-1. Forest and Domain Functional Levels .....	71
Table C-2. Flexible Single-Master Operations Roles .....	77
Table C-3. Automatic Trust Types .....	86
Table C-4. Manual Trust Types .....	87
Table C-5. Anonymous Access Settings.....	89
Table C-6. AD Port\Protocol Use .....	93
Table C-7. Synchronization Port\Protocol Use.....	94
Table C-8. Support Tools Access Permissions .....	97
Table C-9. AD Data Access Permissions .....	104
Table C-10. AD Database Object Access Permissions.....	105
Table C-11. Domain Partition Object Audit Settings .....	108
Table E-1. Cipher Attribute Values .....	132
Table E-2. Server Software Directories .....	133
Table E-3. Administration Server Software Directories and Files .....	133
Table E-4. Server Instance Software Directories and Files .....	133
Table E-5. Password Policy Controls – Content, History, Age .....	134
Table E-6. Server Instance Data Directories and Files .....	137
Table E-7. Directory Database Object Access.....	139
Table E-8. Access Log Controls .....	142
Table E-9. Audit Log Controls .....	143
Table E-10. Error Log Controls .....	143
Table E-11. Password Policy Controls – Lockout .....	146

## 1. INTRODUCTION

This Directory Services Security Technical Implementation Guide (STIG) provides security configuration guidance for the implementation of directory services deployed within the Department of Defense (DoD). The document specifies general requirements applicable to directory service software and specific requirements for Microsoft Active Directory (AD), Microsoft Active Directory Application Mode (ADAM) (to a very limited extent), and Red Hat Directory Server (RHDS). This STIG also provides general guidance for synchronization products that might be used in conjunction with directory servers.

This document replaces the Active Directory STIG, version 1 release 1, dated 10 March 2006. Some requirements from that document have been rewritten in generic terms, but there are no substantial changes to the requirements previously identified for AD.

In this document the term directory service refers to components that collectively provide an information service that makes current and secure directory data available to clients. Although directory server software is the prime element, the directory database, replication mechanisms, and administrative tools all have to be considered directory service components.

Although the data maintained by a directory service can vary, most directories are designed to provide a distributed repository for identification and authentication data. The need to secure this data is clear. As a practical matter of interoperability, all current directory services have to provide some level of support for the Lightweight Directory Access Protocol (LDAP) standard so that various types of clients can more easily access directory data.

This document must be used in conjunction with the other STIGs developed by the Defense Information Systems Agency (DISA). The *Windows 2003/XP/2000/Vista Addendum* and the associated Microsoft documents provide crucial guidance for securing the Windows operating system (OS) on which AD and AD synchronization products execute. The UNIX STIG provides comparable guidance for securing various versions of UNIX and Linux on which Red Hat Directory Server and other directory server software execute. The STIGs that cover database and web server products provide guidance to ensure those services used by some directory synchronization products also support a secure environment.

Other documents that should be used in conjunction with this STIG are the *Active Directory User Object Attributes Specification*, the *Concept of Operations for Global Information Grid Enterprise Active Directory*, *DoD Instruction 8551.1, Ports, Protocols, and Services Management (PPSM)*, and the technical bulletins that are published by the Joint Task Force - Global Network Operations (JTF-GNO).

The *Active Directory User Object Attributes Specification* was developed to provide common naming and attribute guidance to DoD Components that deploy AD. The document specifies the naming convention and acceptable values for some of the attributes of AD User objects. Compliance with the specification is mandated by DoD Chief Information Officer (CIO) policy and is important in supporting interoperability between the Component AD deployments.

The *Concept of Operations for Global Information Grid Enterprise Active Directory* states, “The purpose of this CONOPS is to describe the current state of AD throughout DoD, the technical elements that comprise AD, the current DoD policy and guidance that governs and assists DoD Components executing AD, and the challenges and specifics of how DoD leadership plans to manage current and future AD implementations across the GIG in the NetOps environment.” Components are directed to that document for DoD-specific implementation direction.

*DoD Instruction 8551.1, Ports, Protocols, and Services Management (PPSM)* impacts network implementation. The hosts on which directory server and directory synchronization software execute utilize several network services. Because vulnerabilities have been documented for some of these services, DoDI 8551.1 defines restrictions on the use of the associated ports, protocols, and services in order to protect network-accessible DoD resources. Requirements that reference DoDI 8551.1 are defined in this document.

A final important reference is to the documents published by the JTF-GNO. Bulletins have been written to describe AD and the related DoD initiatives, the operational impact of an AD compromise, and issues involved in implementing Virtual Private Network (VPN) solutions.

## **1.1 Background**

Directory services have been used for a number of years to manage resource metadata that is useful to many clients. As the client/server application model became prevalent, the need for a directory of resources became more important. The advantage of directories may have first become obvious in implementations of directory-enabled applications. However, the current integration of directory service with network operating system (NOS) has provided unmistakable evidence of how valuable a scalable directory service can be when used to enable resource sharing on a large scale.

Early directory service standards were formally defined in the X.500 specifications. However, those standards were classified by some as “cumbersome” and the succeeding LDAP standards eventually became more widely implemented. The availability of industry standards and an open source implementation (OpenLDAP) helped to motivate the development of commercial products that provided LDAP-compatible directory services.

Microsoft introduced AD in Windows Server 2000. Although some of the function and terminology was present in the Windows NT Server OS, the implementation of AD in Windows Server 2000 marked the real incarnation of a native directory service from Microsoft. As such, AD is able to act as a NOS-integrated data store for multiple kinds of enterprise information.

As it is tightly integrated with Windows, AD is less suitable for use with directory-enabled applications. AD is the repository of Windows identification and authorization data, so there is sufficient reason to restrict updates to only trusted administrators. To address this consideration, Microsoft introduced ADAM. By creating a directory service that did not run under highly-privileged OS credentials and that allowed multiple instances per host, Microsoft created a more appropriate directory service for applications.



In the late 1990's the Netscape Corporation developed the Netscape Directory Server product. As the result of company acquisitions, alliances, and product divestitures, a successor version of this product is currently marketed as RHDS. It is not tightly integrated with any NOS, and can be adopted for use to support directory-enabled applications or as an authorization database for OS implementations.

The importance of these individual technologies to DoD is discussed later, but it is appropriate to say that the secure operation of directory services is critical for DoD Components. Whether a directory service provides NOS-integrated access control, a service for directory-enabled applications, or a simple white pages service, the confidentiality, integrity, and availability of directory data must be maintained for everyday operations throughout DoD.

## 1.2 Authority

DoD Directive 8500.1 requires that "all IA and IA-enabled IT products incorporated into DoD information systems shall be configured in accordance with DoD-approved security configuration guidelines" and tasks DISA to "develop and provide security configuration guidance for IA and IA-enabled IT products in coordination with Director, NSA." This document is provided under the authority of DoD Directive 8500.1.

The use of the principles and guidelines in this STIG will provide an environment that meets or exceeds the security requirements of DoD systems operating at the Mission Assurance Category (MAC) II Sensitive level, containing sensitive information.

The Information Operations Condition (INFOCON) for the DoD recommends actions during periods when a heightened defensive posture is required to protect DoD computer networks from attack. The IAO will ensure compliance with the security requirements of the current INFOCON level and will modify security requirements to comply with this guidance.

The JTF-GNO has also established requirements (i.e., timelines), for training, verification, installation, and progress reporting. These guidelines can be found on their web site: <https://www.jtfgno.mil>. Initially, these directives are discussed and released as Warning Orders (WARNORDs) and feedback to the JTF-GNO is encouraged. The JTF-GNO may then upgrade these orders to directives; they are then called Communication Tasking Orders (CTOs). It is each organization's responsibility to take action by complying with the CTOs and reporting compliance via their respective Computer Network Defense Service Provider (CNDSP).

**NOTE:** Field Security Operations (FSO) support for the STIGs, Checklists, and Tools is only available to DoD Customers.

## 1.3 Scope

This document describes security requirements to be applied to implementations of directory services in DoD environments. The information is designed to assist Security Managers, Information Assurance Managers (IAMs), Information Assurance Officers (IAOs), and System Administrators (SAs) with the implementation of more secure directory service configurations. As noted in the previous section, application of the requirements is intended to provide a certain level of assurance. Individual sites must determine the level of assurance that is appropriate to their environment and mission.

This document provides general security guidance for directory server products and for vendor or locally developed solutions that perform directory synchronization functions. Specific security guidance is provided for AD as implemented on computers running the Windows 2000 Server or Windows Server 2003 OS and for RHDS implemented on computers running Linux or UNIX.

## 1.4 Writing Conventions

Throughout this document, statements are written using words such as “**will**” and “**should**.” The following paragraphs are intended to clarify how these STIG statements are to be interpreted.

A reference that uses “**will**” indicates mandatory compliance. All requirements of this kind will also be documented in the italicized policy statements in bullet format, which follow the topic paragraph. This makes all “**will**” statements easier to locate and interpret from the context of the topic. The IAO will adhere to the instruction as written.

Each policy bullet includes the STIG Identifier (STIGID) in parentheses that precedes the policy text and references the corresponding vulnerability check in the SRR Checklist and Vulnerability Management System (VMS). An example of this will be as follows: “(*G111: CAT II*).” If the item presently does not have a STIGID, or the STIGID is being developed, it will contain a preliminary severity code and “N/A” (i.e., “[*N/A: CAT III*]”). Throughout the document accountability is directed to the IAO to “ensure” a task is carried out or monitored. These tasks may be carried out by the IAO or delegated to someone else as a responsibility or duty.

A reference to “**should**” indicates a recommendation that further enhances the security posture of the site. These recommended actions will be documented in the text paragraphs but not in the italicized policy bullets. All reasonable attempts to meet this criterion will be made.

## 1.5 Vulnerability Severity Code Definitions

During a Security Readiness Review, reports are produced that detail the vulnerabilities that are the result of deviations from the STIG requirements. These vulnerabilities are classified by severity into the following categories:

<b>Category I</b>	<p>Vulnerabilities that allow an attacker immediate access into a machine, allow superuser access, or bypass a firewall.</p> <ul style="list-style-type: none"><li>- Category I vulnerabilities include those that allow an attacker immediate access to authentication data in the directory or to objects in the directory that are used to control access on other systems. Storing unencrypted passwords in a directory or enabling anonymous access to non-public data are two examples. An unauthorized Windows trust relationship between AD domains is another example.</li></ul>
<b>Category II</b>	<p>Vulnerabilities that provide information that has a high potential of giving access to an intruder.</p> <ul style="list-style-type: none"><li>- Category II vulnerabilities include those that result in the disclosure of information that has a high potential to allow an attacker to gain access into a directory or to objects in the directory that are used to control access on other systems. Allowing everyone read access to all directory database account data is an example.</li></ul>
<b>Category III</b>	<p>Vulnerabilities that provide information that potentially could lead to compromise.</p> <ul style="list-style-type: none"><li>- Category III vulnerabilities include those that result in the exposure or loss of information that could lead to the compromise of a directory or to objects in the directory that are used to control access on other systems. One example would be the use of weak data signing algorithms that allow an attacker to intercept and modify directory data as it traverses a network. Another example would be the failure to maintain adequate documentation needed to efficiently restore a directory service.</li></ul>

**Table 1.1. Vulnerability Severity Code Definitions**

## **1.6 DISA Information Assurance Vulnerability Management (IAVM)**

The DoD has mandated that all IAVMs are received and acted on by all commands, agencies, and organizations within the DoD. The IAVM process provides notification of these vulnerabilities and alerts require that each of these organizations take appropriate actions in accordance with the issued alert. IAVM notifications can be accessed at the JTF-GNO web site: <https://www.jtfgno.mil>.

## **1.7 STIG Distribution**

Parties within the DoD and Federal Government's computing environments can obtain the applicable STIG from the Information Assurance Support Environment (IASE) web site. This site contains the latest copies of any STIG, as well as checklists, scripts, and other related security information. The NIPRNet URL for the IASE site is <http://iase.disa.mil/>.

## **1.8 Document Revisions**

Comments or proposed revisions to this document should be sent via e-mail to [fso\\_spt@disa.mil](mailto:fso_spt@disa.mil). DISA FSO will coordinate all change requests with the relevant DoD organizations before inclusion in this document.

## **2. DIRECTORY SERVICE ELEMENTS**

### **2.1. Introduction**

As it is used in this document, the term directory service refers to the components that collectively provide an information service that makes current and secure directory data available to clients. This description is important to a security perspective because all of the elements in a directory service have to be properly configured and protected if there is to be an expectation that the confidentiality, integrity, and availability of directory data can be maintained.

In the following discussions, the elements of a directory service are roughly divided into directory server software, the directory database, replication mechanisms, and administrative tools. This is not meant as a formal itemization, just a means of dividing a directory service into more manageable parts for discussion. While vendor implementations are not aligned simply with these elements, each implementation has to address the elements in order to provide a directory service. The discussions in this section are intended to provide background on directory service security issues. Section 3, Directory Service Security Requirements, provides the specific requirements that must be met in order to address the security issues.

LDAP is the industry standard protocol for access to directory services. While it is not a component of a directory service, support for current LDAP standards is a practical requirement and a security essential for commercial directory service software. LDAP has been an evolving standard with current activity focusing on LDAP version 3 (LDAPv3). IT industry groups recognized that the LDAP version 2 standard did not incorporate sufficient authentication controls and should not be used. Work within the Internet Engineering Task Force (IETF) has resulted in a number of documents that address LDAP security issues. One of the latest is *Request For Comment (RFC) 4513, Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms*. RFC 4513 addresses required security support including mechanisms for authentication and data integrity and confidentiality. Along with related LDAP RFCs, RFC 4513 obsoletes the earlier LDAP version 3 base RFC (2251) and the security documents, RFCs 2829 and 2830.

Those interested in background information on directory services and LDAP should review a primary reference on the subject, *Understanding and Deploying LDAP Directory Services, Second Edition*, by Howes, Smith, and Gordon.

### **2.2. Common Directory Service Elements**

As stated earlier, each directory service provides server software, a directory database, replication mechanisms, and administrative tools. The following subsections address each of these elements with respect to the high-level security considerations for them. In addition, there is a brief discussion of the network ports and protocols that are associated with LDAP-based directory services.

### **2.2.1. Directory Server**

The primary element in every directory service solution is the directory server software. It is responsible for a number of essential functions:

- Providing a central point for clients to access directory data
- Ensuring that access to directory data is in accordance with the desired security policy
- Enabling secure access to clients across networks that may not have adequate confidentiality or integrity features
- Ensuring that data integrity issues are addressed when data is updated
- Enabling distribution of directory data where needed to address availability and performance issues
- Enabling coherent backup of directory data in such a way that data availability is maintained and that data can be restored when necessary.

The manner in which these functions are accomplished can raise many security considerations. In the following sections, authentication, account definitions, and query and update access are specifically addressed. Following those discussions, several general server considerations are described.

#### **2.2.1.1. Authentication**

Because authentication verifies the identity of a user, it is crucial for controlling access to resources and collecting meaningful audit data. Directory services within DoD are likely to hold data for applications or systems at the sensitive or classified confidentiality level as defined in DoD Instruction 8500.2. But even if the confidentiality level of the data is public and minimal access control is necessary, authentication mechanisms must still be available to distinguish administrative versus non-privileged access to the directory.

The IETF RFC 4513 standard discusses authentication as it relates to LDAPv3-based directory servers. During the process known as the bind operation (or binding), a user may be authenticated by a directory server. The standard describes different methods of authentication:

- The simple authentication method includes three authentication mechanisms: anonymous, unauthenticated, and name/password. The anonymous and unauthenticated mechanisms amount to unauthenticated access because there is insufficient data to verify a client's identity.
- The Simple Authentication and Security Layer (SASL) method refers to the standard (RFC 4422) SASL framework for authentication and data security services. SASL defines several authentication mechanisms including ANONYMOUS, PLAIN, DIGEST-MD5, EXTERNAL, and GSSAPI. ANONYMOUS and PLAIN provide similar processing to the simple anonymous and name\password mechanisms. DIGEST-MD5 uses challenge/response processing with a shared secret. The GSSAPI mechanism uses Kerberos and EXTERNAL uses network services such as Secure Sockets Layer / Transport Layer Security (SSL/TLS).

The following security considerations apply to these authentication mechanisms:

- Any mechanism that does not result in the authentication of a specific account is generally not useful to DoD because it amounts to anonymous access. Directories at the public confidentiality level are the only case where anonymous access may be appropriate.
- Mechanisms such as simple name\password and SASL PLAIN allow passwords to travel between client and server in clear text. This is unacceptable over a network unless accompanied by encryption of the session such as with SSL/TLS.

The following considerations apply to features that a directory server may support in association with authentication processing:

- The use of SSL/TLS in a Public Key (PK)-enabled configuration is incomplete without some form of certificate validation. Integration of Certificate Revocation List (CRL) or Online Certificate Status Protocol (OCSP) processing is needed to ensure revoked certificates cannot be used.
- Limits on the source location of the client, implemented through IP address or host name checks, have limited application. If the address or host name is supplied by components of an un-trusted network, their accuracy cannot be considered reliable.

#### **2.2.1.2. Account Definitions**

Directories typically hold client account definitions that include identification and authentication data. Although the use of digital certificates is becoming widespread, user IDs and passwords have been the most common forms of identification and authentication data stored in account definitions in directories. The obvious importance dictates that controls are configured to enforce the use of stronger forms of authentication data.

To ensure the use of strong passwords, DoD security policy on password characteristics must be enforced for passwords stored in directories. To deter password cracking attacks, there must be facilities that provide automated account lockout after repeated, unsuccessful logon (bind) attempts. When digital certificates are used, there must be facilities to address expiration and source (certificate authority) concerns.

Components of directory server software or application code might be used to implement the required controls. The use of configurable options in directory server software is generally a preferred solution. This is because it usually keeps the IA functions isolated from application code, provides a central point of control, and allows easier configuration changes. However, the evaluation of multiple server configurations and directory product capabilities can lead to the conclusion that security controls are best implemented in the application.

The following considerations apply to account definition and lockout:

- Controls for password complexity, including length and composition, are required.
- Controls for password history are required to ensure passwords are not frequently reused.

- Controls for the expiration of authentication data are required. This includes expiring passwords and checking for expired and revoked digital certificates. Considerations for directory service process accounts are needed to reduce operational impacts.
- Controls are needed to ensure repeated unsuccessful logon attempts within a brief time period result in a lockout of the account.
- The use of a formally managed Public Key Infrastructure (PKI) is needed to ensure interoperability and adherence to strong credential assignment policies. [Self-signed certificates inherently reduce interoperability.]

### 2.2.1.3. Query and Update Access

Support for client access for query and update is the primary reason for the creation of a directory service. Directory servers provide the central connection point for a client to access directory data. To preserve the confidentiality, integrity, and availability of directory data, server software has to ensure query and update access is secure.

The following considerations apply to general directory server support for query and update access:

- Support for authentication of the client by the server is necessary to establish the identity under which access is to occur. Without authentication, access control of directory data is impossible. Authentication is discussed above, but it is mentioned again to emphasize its importance to directory service security.
- Support for authentication of the server by the client is necessary in some circumstances. In the case of administrative actions and replication between servers, it is essential that the client is able to establish that the server represents the intended directory service.
- Support to protect the integrity of the data in transit over a network between server and client is sometimes necessary. This support can be from either the directory server or underlying network services. Data signing is usually implemented to meet this need.
- Support to protect the confidentiality of data in transit over a network between server and client is sometimes necessary. This support can be from either the directory server or underlying network services. This means formally validated encryption must be implemented.

To enable related directory data to be accessible when it is distributed across multiple servers, the concept of a knowledge reference was devised. Knowledge references are ways to allow directory information to be retrieved when that information is not within the directory that was queried. There are two general types of knowledge reference:

- A referral is an LDAP knowledge reference that is provided by the server to the client to inform the client which directory could be queried to obtain the desired data. The client is responsible for performing additional queries.
- Chaining is an LDAP knowledge reference that is used by a server to contact another server to obtain the directory data sought by a client. The server contacted by the client returns the results passed from the server that holds the data, eliminating the need for the client to perform queries on additional servers.



The following security considerations apply to knowledge references:

- Control over the definition of referrals is important because they could become a form of directory poisoning, sending clients to non-existent servers or servers with invalid data.
- Authentication issues arise when chaining is performed. This is because the server on which the data resides is actually being queried by another server, not the client. Typically some form of proxy authentication is used, but careful configuration is required to ensure data is not returned to unauthorized clients.
- For either form of knowledge reference, local consideration must be given as to whether references to directories outside the control of the originating organization are desirable. A directory hosted by an outside organization may not be subject to the same security controls. As a result, there may be less confidence in the integrity of the data.

Directory Services Markup Language (DSML) is a standard designed to allow the exchange of directory data in Extensible Markup Language (XML) format. DSML enables web service access by supporting Simple Object Access Protocol (SOAP) over Hypertext Transport Protocol (HTTP) or Hypertext Transfer Protocol over SSL (HTTPS). DSML is commonly implemented in web server-based gateways that receive DSML-format data from clients and pass the data in LDAP format to a directory server.

The chief security consideration for a DSML implementation is related to authentication. Similar to the issue with directory chaining, care is needed to ensure the identity used by the DSML gateway to access the directory server does not allow inappropriate access to data.

#### **2.2.1.4. Other Server Issues**

Beyond the authentication, account definition, and query and update access issues, there are a number of other security issues that are associated with the directory server. Some of these are common to all types of server applications; others have particular relevance to directory servers.

The following security considerations are general application issues that apply to directory server software:

- The integrity of the server software executables and configuration files must be maintained. Proper file access control is needed to prevent inadvertent or malicious updates that compromise server software integrity.
- Unplanned and uncoordinated changes to software and configurations can affect server availability. A documented configuration management process addresses these issues.
- If locally written programs are used to supplement or enhance the function of a directory server, it is important that the code is subject to configuration control and is factored into disaster recovery plans.
- The OS identity under which a directory server instance executes has to be considered. While some configurations require the use of highly privileged credentials such as the UNIX superuser, not all configurations do. In those cases where highly privileged credentials are used, it is more important that other functions such as gateways are not executed in the same server instance.

The following security considerations have unique importance to directory servers:

- Because directory data is sometimes used as the basis on which access control decisions are made, a server failure could result in a loss of availability for the platforms that rely on the directory. To mitigate this vulnerability, it may be necessary to deploy redundant directory servers to ensure continuous availability of directory data.
- If directory data includes authentication data such as passwords that can be reused if obtained, the server must protect the confidentiality of that data in the database. This means formally validated encryption must be implemented.
- Directory architecture implementations that involve fractional or selective replication may require special restoration sequences in the event of failure or loss of multiple servers. A directory server architecture diagram (or text documentation) can provide key information needed for recovery.
- Directory server software that allows updates to be performed on multiple servers must support update conflict resolution. One data element that may be used in conflict resolution is transaction time. Time of day is also be used in some authentication mechanisms to deter replay attacks. For these reasons, a form of time synchronization with an authorized time source is essential for each directory server.

### **2.2.2. Directory Database**

As the directory server is the central access point for directory data, the directory database is the central repository of the data. Although there are exceptions, directory databases are commonly implemented using specialized database engines rather than conventional relational database management software. This reflects the fact that the usual access pattern for directory databases involves a relatively large number of queries, compared to a small number of updates. As a consequence of this implementation, there are some security concerns that are similar to those for other database systems and some that are unique.

Two characteristics of directory databases influence security design. The first is the fact that directory data is represented as entries known as objects and object attributes. The second is that LDAP-based directories are logically structured in a tree form. That is, there is a root from which branches extend. Entries reside figuratively at the root, at branch points, and as leaves.

There is a special directory entry accessible on each directory server at the root of LDAPv3 directories; it is known as the root Directory System Agent (DSA) Specific Entry (DSE). The root DSE holds several important pieces of information, expressed as attributes. The information includes the versions of LDAP (v2 or v3) supported, information on supported security mechanisms, and information on the naming contexts (major directory tree branches or subtrees) that are held in the directory.

There are some unique considerations for controlling access to the root DSE:

- Some clients attempt to use anonymous access while negotiating security mechanisms. In this case, the client may query the directory to obtain the supported mechanisms specified in the root DSE before attempting to bind with the actual client credentials.

- Clients using a commercial-off-the-shelf (COTS) product for query, update, or monitoring might attempt to query the root DSE to determine which LDAP options are supported before attempting to use the options. As with the security mechanisms, this root DSE query might be attempted before the bind with the client's credentials.

In addition to the root DSE, access control over all the objects and attributes in the directory database is required in order to preserve the confidentiality, integrity, and availability of the data. Although control mechanisms vary, a form of access control inheritance that flows through the directory tree branches is desirable. Implementations may refer to the statements that control access as Access Control Lists (ACLs) or Access Control Instructions (ACIs).

Anonymous or unauthenticated (which is considered equivalent in this document) access to directory database objects is a significant security consideration. Generally, the decision whether to permit anonymous access depends on the nature of the data. It is possible to configure a working directory that does not permit anonymous access to any data. This is desirable within DoD because the directories commonly contain data that is itself designated as sensitive or classified, or data that might permit access to other systems that process sensitive or classified data. However, there may be cases where the client tools have been written to use some anonymous access as indicated in the earlier discussion on the root DSE.

The following security considerations are general database issues that apply to directory databases:

- Databases are typically composed of multiple files that reside in an OS-supported file system. As with any file system data, proper file access control is needed. Appropriate write permissions prevent inadvertent or malicious updates that could compromise database integrity. Restrictive read permissions prevent the capture of the data so that it cannot be subjected to offline attacks.
- There must be support within database object permissions or through server facilities to capture audit data. Without a way to determine the source and extent of an access, it may be impossible to determine the nature or extent of damage if the directory server is compromised.
- A database-specific backup and restore capability is usually an essential requirement. Because databases are composed of multiple files and coordinated updates are critical to data integrity, the backup and restore solution must often be specific to the database.
- Depending on database design, there may be referential integrity issues. In relational database terms, this means that updates in one table may need to be coordinated with updates in another. The database or server software must provide facilities to address any referential integrity issue.

A final area of consideration for directory databases is protection of the directory schema. As with other databases, the schema provides the rules to which data must conform. Clients and servers can consult the directory schema to understand the permitted format and content (e.g., size and range) of specific objects and attributes. LDAPv3 directory databases contain their directory schema as LDAP attributes that can be queried.

The need to protect the directory schema is apparent. A compromise of the directory schema could literally lead to the loss of integrity of directory data. Deletions of objects and attributes defined in the schema can cause queries for directory entries to fail. If queries for authentication or authorization data fail, access control decisions on systems using that data could result in denying authorized access or allowing unauthorized access.

The following considerations apply to schema security:

- Because the directory schema is composed of entries in the directory database, the normal access control mechanisms for database entries must be properly configured. In general, using access control specified at higher levels where the schema is defined, with inheritance to the lower levels, is the proper approach.
- Unplanned and uncoordinated changes to directory schema can affect server availability and data integrity. A documented configuration management process addresses these issues.

Although it is not explicitly a security issue, the use of a standard schema (or a schema with common core elements) by all organizations that may need to exchange directory data is a significant interoperability issue. The *Concept of Operations for Global Information Grid Enterprise Active Directory* requires that information be collected to document changes made to AD schema implementations by DoD components. Please refer to that document for specific requirements.

### **2.2.3. Replication**

Directory data replication refers to the process in which data is copied from one directory server to another. In general, replication is intended to enhance directory availability and performance by allowing clients to access the same data from more than one directory server. It makes the data accessible to a higher number of directory clients and can better support clients over a wide geographic area. It can be part of a strategy to recover from network or individual server outages. It may also be used in the restoration process for a failed directory server.

Various replication configurations are possible:

- Multi-master replication refers to a configuration in which multiple servers may accept updates to copies of the same data. The servers participating in multi-master replication exchange data with each other and a conflict resolution process addresses updates to the same directory entries.
- In single-master replication, one server maintains a read-write copy of the directory and changes are propagated to one or more servers that hold read-only copies.
- Fractional replication refers to replication in which some directory entries are excluded. This is useful for performance or security reasons where all the data is not wanted on every server.
- Implementations might use a combination of multi-master and single-master replication for a single directory. An example of this is seen in AD where schema data is replicated using a single-master configuration and account data is replicated using multi-master replication.

The following security considerations apply to the use of replication:

- Servers participating in replication have to perform mutual authentication. That is, each server must identify and authenticate the other. This ensures that data is sent from an authorized source to an authorized target.
- Support to protect the confidentiality of replication data in transit over a network between servers is necessary in some cases. This support can be from either the directory server or underlying network services. This means formally validated encryption must be implemented. Data in transit must be protected when it contains authentication data such as passwords that can be reused if obtained. In addition, if a sufficient aggregate of unencrypted directory data is intercepted, that aggregate could constitute sensitive data.
- Support to protect the integrity of replication data in transit over a network between servers is necessary. This support can be from either the directory server or underlying network services. Data signing is usually implemented to meet this need.
- Replication schedule parameters are important to maintaining current data. For environments in which multiple directory servers provide data used in identification, authentication, or authorization decisions, timely replication is critical to enforcing current access policy. When replication schedules do not ensure current data is available, access could be incorrectly permitted or denied.
- Some implementations allow access control for some entries to be based on attributes of other entries. If fractional replication is used, the data needed for access control decisions could be missing in the replicated copy. Careful design of fractional replication specifications is needed to avoid this issue.

#### **2.2.4. Administration Tools and Accounts**

Given the standard practice of performing administrative tasks from workstations rather than at server consoles, the security of administration tools and accounts is a common consideration for server applications. This reflects the fact that administrative authentication data as well as potentially sensitive directory data is flowing over one or more networks.

The following security considerations apply to the use of administrative tools:

- The confidentiality of any transmission of administrator authentication data in plain-text form has to be protected.
- The confidentiality of any transmission of authentication data in plain-text form that is being read from or written to directory entries must be protected.
- If any of the directory data is designated classified, or if the data could be used to access systems that process classified data, the confidentiality of that data has to be protected in transmissions across all DoD networks.
- If any of the directory data is designated sensitive, or if the data could be used to access systems that process sensitive data, the confidentiality of that data has to be protected in transmissions across wireless or non-DoD networks.
- When administrative tasks are implemented through components such as web servers that could be used for other tasks, the components associated with the administrative

functions should be dedicated to those tasks. This reduces the risk that a compromise related to non-privileged use could also affect the administrative tasks.

Many commercial and open-source LDAP clients support the use of SSL/TLS to provide encryption as the means to address the data transmission security considerations.

As with administrative tasks for other server applications, there are concerns about the privileged accounts that are used for directory administration. The following considerations apply:

- Passwords used with administrative accounts must at least conform to the minimum standards for all passwords, but should be subject to more restrictive password policies (especially greater length). There may also be circumstances, such as when heightened threat levels exist, that more restrictive policies are required on an interim basis. The use of digital certificates for authentication is preferred.
- Accounts with specific privileges should be used to perform administrative tasks instead of accounts with a broad range of high privileges. As one example, if an account can be specified for use by replication, that account should be used exclusively for replication processes.
- The number of accounts with administrative privileges has to be limited.
- Documentation must exist to detail accounts to be used for administrative tasks. Periodic comparisons of the documentation to system definitions can be used for validation.
- When possible, the vendor standard account names for administrative accounts should be replaced.

### **2.2.5. Network Ports and Protocols**

While network ports and protocols are not actual elements of a directory service, they must be enabled within the underlying host and network infrastructure for a directory service to function. Directory servers are both network service providers and consumers. They always act as providers by listening for connections from clients and providing responses. They can act as consumers when they participate in replication and contact other servers to propagate directory updates.

The following are the standard directory service network ports and protocols:

- Port 389 is the standard port designated by the Internet Assigned Numbers Authority (IANA) for LDAP communication.
- Port 636 is the standard port designated by the IANA for Lightweight Directory Access Protocol over SSL (LDAPS) communication.

Vendor implementations usually allow ports other than the IANA standards to be configured for LDAP or LDAPS. This is used to allow multiple directory servers to be active on a single host or to avoid OS restrictions on non-privileged use of the IANA Well Known Ports (0-1023).

There are other ports (some unique) associated with the different vendor directory server implementations. Typically these other ports are related to administrative tools, security services

(such as Kerberos), or functions of a NOS-integrated directory service. Please refer to the technology specific appendices in this document for discussions of these other ports.

The following security considerations apply to network ports and protocols:

- Traffic over port 389 may, or may not, be encrypted. It depends on the vendor implementation and site customization. [Current AD implementations, configured securely, do encrypt administrative traffic on port 389.] When encryption is not enabled, plain-text authentication data, or any other sensitive data, could be intercepted in transit.
- The Start TLS operation can be used by implementations using port 389 to provide selective encryption of session data.
- Traffic over port 636 uses the SSL/TLS protocol. Properly configured, this provides confidentiality and integrity protections to the session.

DoD policy specifies restrictions on where these ports may be enabled. These restrictions are described in *DoD Instruction 8551.1, Ports, Protocols, and Services Management (PPSM)* and the associated *Ports, Protocols, and Services (PPS) Assurance Category Assignments List*. The *PPS Assurance Category Assignments List* provides the detailed guidance for specific ports and protocols. It is available to DoD and Government users through the IASE web site.

In order to support directory service environments in which LDAP traffic is transmitted across DoD enclave boundaries, but remain compliant with the DoD policy, the use of encrypting VPN technology is the only known acceptable solution. Please refer to the *Network Infrastructure STIG* for information on VPN implementation.

### **2.3. Directory Synchronization Tools and Technology**

As in other large organizations, the implementation of directory services within DoD has not been uniform. This is the predictable result of varying procurement cycles, budget challenges, and mission needs. Even when individual deployments followed similar timelines, isolated directory service environments have been built to provide necessary security boundaries.

In the case of Microsoft Windows environments, there has been a proliferation of isolated directories. In specific terms, this means that more AD forests exist than are required in a purely technical sense. In research compiled by the DoD Active Directory Interoperability Working Group (DADIWG) in 2005, there were indications that more than 1,500 AD forests existed within DoD.

To support DoD's net-centric strategy, Components have had to implement solutions that allow information in multiple directories to be shared and synchronized. Although some software is bundled with Windows and UNIX OSs to support this, the need for more robust solutions has resulted in the development of COTS synchronization products. This section provides a very brief description of a few of the Windows components and COTS products that support synchronization of AD data.

Windows includes two tools that support the bulk import and export of AD data:

- The Comma Separated Value (CSV) Data Exchange (CSVDE) program provides import and export of AD data in comma-separated value format. CSVDE contacts the AD domain controller on port 389 or a Global Catalog server on port 3268.
- The LDAP Data Interchange Format (LDIF) Data Exchange (LDIFDE) program provides import, export, and modification of AD data using LDIF files. LDIF is a file format standard for directory service data. LDIFDE can contact a domain controller on port 389 or a Global Catalog server on port 3268. If a domain controller is configured to support SSL, LDIFDE can use ports 636 or 3269 and utilize SSL-based encryption. Later versions of LDIFDE also include a switch ("-h") that invokes SASL-based encryption.

The primary security issues with CSVDE and LDIFDE are protection of the data files. Whether in transit or in a file system, the data could contain sensitive information or the aggregate of the information could be sensitive. Unless an encrypting mechanism is explicitly specified, the data would be vulnerable to disclosure if intercepted in transit.

CPS Systems sells the SimpleSync COTS product. SimpleSync performs directory synchronization operations between several different types of directory services including AD and other LDAP-based products. SimpleSync can communicate with domain controllers over LDAP on port 389 or LDAPS on port 636.

Microsoft sells the Microsoft Identity Integration Server (MIIS) as a COTS directory tool. Microsoft offers the Identity Integration Feature Pack (IIFP) as a downloadable feature for Windows Server 2003. IIFP supports a subset of the directory services supported by MIIS, but otherwise both components provide the same type of service. This service is able to gather and synchronize directory data from a number of sources including AD. MIIS stores its data in a Microsoft SQL Server database and allows complex manipulation of the directory information. MIIS is the successor to the Microsoft Metadirectory Services (MMS) product that was available for Windows 2000 Server. Microsoft plans to offer Microsoft Identity Lifecycle Manager (ILM) 2007 as a product that includes and enhances the functionality of MIIS.

DSML Services for Windows (DSfW) is a downloadable feature from Microsoft. DSfW enables access to AD by using SOAP over HTTP or HTTPS. DSfW requires the use of a gateway program running under the Microsoft Internet Information Services (IIS) web server. The Directory Services Data Exchange (DSDE) program is available with DSfW to use as a client to access DSfW.

Without respect to the specific product or technology used, the following general security issues must be considered for directory synchronization tools and technology:

- Data files holding substantial aggregations of directory data can become sensitive. Examination of a sufficient aggregate might disclose force strength or composition.
- The exchange of "contact object" data such as that used in Global Address List (GAL) synchronization does not normally represent sensitive traffic because that data is not used in identification, authentication, or authorization operations.



- Permitting anonymous access to directory data can only be considered acceptable if the confidentiality category of the associated system is public.
- Mutual authentication, that is authentication of the client by the server and the server by the client, is necessary to help deter a spoofing attack. Invalid directory data could result in the disclosure of sensitive data if that directory data is used in access control decisions.
- The use of query and update restrictions such as quotas can help to deter denial of service attacks.
- The vulnerability of certain network services must be taken into account when considering access to directory services in a network that spans enclave boundaries.
- Limiting write access to the directory through access permissions, quotas, or other update controls can help to prevent inadvertent or erroneous overwrite of directory data or flooding of the target directory.

This page is intentionally left blank.

### 3. DIRECTORY SERVICE SECURITY REQUIREMENTS

Directory services are commonly responsible for managing and providing access to critical organization data. This is unquestionably true of directory services in which identification, authentication, and authorization data is stored for reference by OS or application security components. For NOS-integrated directory services such as AD, there is a very close relationship between directory and client security. If the confidentiality, integrity, or availability of the related directory service is compromised, it is likely that the security of a dependent server or workstation is also compromised.

The threats to directory service components have much in common with the threats to other software and data components. But because a large number of hosts may depend on the data, the vulnerabilities can be more serious. Users may disrupt availability through malicious or simply unintended operations on the directory. When a user with administrative-level privileges initiates a harmful action, the consequences can range from a small-scale denial of service to the effective disruption of all the dependent servers and clients.

This section provides the specific security requirements that apply to directory service components and to products that may be used to manipulate directory data. This section is broken into subsections that align with the IA Controls subject areas defined in DoD Instruction 8500.2, Information Assurance (IA) Implementation. These subject areas are as follows:

- Security Design and Configuration
- Identification and Authentication
- Enclave and Computing Environment
- Enclave Boundary Defense
- Physical and Environmental
- Continuity
- Vulnerability and Incident Management.

Some of these areas are further divided to provide a more cohesive presentation. The Personnel subject area is not included as there are no controls in that subject area that are addressed in a directory service security review.

It is important to understand specific terminology used in this and following sections:

- Synchronization products or solutions - There are products by vendors, and in the public domain, that read and update directory data. It is also possible that Components may write their own applications as part of local solutions to do this. These products and solutions can perform various functions from simple reporting to access control configuration and complex identity provisioning. While specific guidance for these products may be provided in future versions of this document, at present the requirements for these products refer to them as synchronization products or solutions.

It is essential to note that the requirements stated in this document are intended for use in a specific context. That context is the high security (sometimes referred to as specialized security - limited functionality) configuration required for DoD Automated Information System (AIS)

components. Such a configuration is achieved by compliance with the other DoD information assurance guidance. This refers specifically to the *Windows 2003/XP/2000/Vista Addendum*, the *UNIX STIG*, the *Domain Name System STIG*, the *Enclave STIG*, and the *Network Infrastructure STIG*. Some AD synchronization solutions utilize COTS web servers and database management systems. In those cases, compliance with the *Database STIG* and the *Web Server STIG* is also assumed.

### 3.1. Security Design and Configuration

This section describes directory service security requirements based on applicable DoDI 8500.2 IA Controls in the Security Design and Configuration subject area. These requirements address five general areas: product design characteristics, configuration and implementation integrity, network services, software integrity, and security service partitioning.

#### 3.1.1. Product Design

Software must be properly designed and implemented to maintain the security of the data it manipulates. When the data is used in identification, authentication, or authorization services, the software is identified as an IA product. Once identified as such, the software must be formally evaluated to establish objectively that it meets certain design and implementation requirements that address potential vulnerabilities. The currently accepted evaluation criteria are specified in the Protection Profiles of the Common Criteria.

Directory server products are commonly used to maintain identification and authentication information. AD is a prime example of this usage. In any case where a directory server is used in this fashion, a formal evaluation of the product is required.

- *(DS00.1100: CAT III) The IAM will ensure the acquisition of directory server products used to maintain identification, authentication, or authorization data meets the applicable Common Criteria, NIAP, or FIPS evaluation and validation requirements specified in NSTISSP No. 11 and DoDI 8500.2.*

With respect to AD, Windows 2000 Server and Windows Server 2003 have been validated for conformance at the Common Criteria Evaluation Assurance Level (EAL) 4 Augmented. These validations meet the formal evaluation requirements for directory components of Windows as an IA product.

When synchronization products update directory data that is used in identification, authentication, or authorization services, they function as IA products and must also be formally evaluated.

- *(DS05.0100: CAT III) The IAM will ensure the acquisition of directory synchronization products that maintain identification, authentication, or authorization data meets the applicable Common Criteria, NIAP, or FIPS evaluation and validation requirements specified in NSTISSP No. 11 and DoDI 8500.2.*

- *(DS05.0110: CAT III) The IAM will ensure directory synchronization products that maintain identification, authentication, or authorization data from sensitive systems meet the medium robustness requirements defined in DoDI 8500.2 when any of the following is true:*
  - *Directory synchronization data traverses public networks.*
  - *Directory synchronization data resides on systems that are accessible by individuals not authorized to access the information.*

It should be noted that when synchronization products are used to perform e-mail contact or GAL synchronization functions they are not normally identified as IA products. However, if the items synchronized are used by an AIS to perform identification, authentication, or authorization, the synchronization products would be identified as IA products.

### **3.1.2. Configuration and Implementation Integrity**

The software components of a directory service are applications that execute under an OS. Even for a NOS-integrated directory service, there are separate OS components that must be available for the directory service to function. Given this relationship, the importance of securing the OS that is the host of the directory service cannot be overstated. To ensure the OS is secure, the applicable DoD guidance must be applied.

- *(DS00.1110: CAT II) The IAO will ensure the applicable OS STIG is applied to the host on which the directory server runs.*

As noted in Section 2.2.2, Directory Database, the schema for a directory database is actually expressed as attributes in the database for implementations that are compliant with the LDAPv3 standard. Invalid modifications to the schema could render directory servers or dependent applications inoperable. To help ensure schema modifications are appropriately designed and implemented, they must be subject to a configuration management process.

- *(DS00.0100: CAT III) If the directory schema is altered by the addition, change, or deletion of objects, the IAM will ensure a documented configuration management process is used for the implementation of those added, changed, or deleted schema elements.*

It is important to note that certain product installation procedures might attempt to perform schema modifications without explicit notification. While the intent is not malicious, the impact could be negative in the short or long term if an object definition is altered and a conflict is created. SAs responsible for product installations are advised to carefully review product documentation for this issue.

Some directory services enable simplified resource access by allowing the user authentication performed under the scope of one directory to be honored under the scope of a different directory. This is sometimes part of solutions referred to as single sign-on implementations. This can reduce the number of accounts a user must have and the number of times the user must authenticate. Prime examples of such a mechanism are certain types of AD trust relationships.

For the discussion here, this mechanism is referred to generically as cross-directory authentication.

A technical review of whatever specific configuration values enable cross-directory authentication is not sufficient to determine if the intended resource access policy is being enforced. It is necessary to maintain a baseline of information about valid cross-directory authentication arrangements. By comparing the documented baseline to the specific configuration values, improper cross-directory authentication can be exposed.

- *(DS00.1120: CAT III) The IAO will ensure documentation is maintained to describe any valid cross-directory authentication configurations.*

The technology-specific appendices of this document describe the cross-directory authentication mechanisms, and the specific documentation that has to be maintained for each type in order to meet this requirement.

Cryptographic algorithms support data encryption and signing functions to maintain data confidentiality and integrity. Directory service implementations may use OS facilities, implement the algorithms internally in server code, or indirectly utilize standard implementations through protocols such as TLS. Unfortunately some of the implementations include algorithm or key length options that are too weak to be acceptable for use in DoD. If a weak algorithm were used, it might be possible for an attacker to access or intercept signed and encrypted directory data, decipher it, and replace it with modified data.

To eliminate the use of weak algorithms, requirements specify the use of implementations that have been validated by the National Institute of Standards and Technology (NIST) under Federal Information Processing Standards (FIPS) specifications.

- *(DS00.1130: CAT II) SAs will ensure directory service software is configured to use FIPS 140-2 validated encryption, key exchange, digital signature, and hash algorithms for all required data signing or encryption functions.*

**NOTE:** This requirement is met for AD by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum*.

Some directory synchronization software can also be configured to use data encryption and signing functions. A comparable requirement applies to these implementations.

- *(DS05.0120: CAT II) SAs will ensure directory synchronization software is configured to use FIPS 140-2 validated encryption, key exchange, digital signature, and hash algorithms for all required data signing or encryption functions.*

### 3.1.3. Network Services

The nature of a directory service means that directory data is transmitted across networks. In some cases directory implementations span DoD enclave boundaries and directory data is transmitted over wide area networks. There are two network security considerations when DoD enclave boundaries are traversed: directory services and other collocated services.

Depending on the implementation, directory data might be transported by various network services over multiple network ports. While the use of the LDAP protocol over port 389 and the LDAPS protocol over port 636 might be the most common, there are other less obvious uses.

The implementation of AD forests is a good example of diverse network port and service usage. Section C.2.6, Ports and Protocols, lists a number of ports and services used by AD and by synchronization solutions for AD. These ports and services enable queries, updates, and data transfers that support identification, authentication, and authorization as well as AD data replication between domain controllers. Opening these ports in network infrastructure components such as firewalls and routers can make the AD services accessible to attack from hosts that have gained unauthorized access to the network.

Even when the directory service components have been made as secure as possible, consideration must be given to the fact that networks are shared resources. Collocated hosts might include Linux systems running specific applications with associated directory servers such as OpenLDAP or file servers using the Common Internet File System (CIFS) protocol on port 445. Because these other systems may have vulnerabilities that make them susceptible to attack, opening ports on firewalls or routers to support directory services could elevate the risk of compromise of other systems on the same network.

Both directory servers such as Windows domain controllers and other hosts using the same network protocols are protected through compliance with the guidance in *DoDI 8551.1, Ports, Protocols, and Services Management (PPSM)*. Since network traffic for some of the services used by directory services is not permitted across any of the network boundaries defined through DoDI 8551.1, it is necessary to employ DoD-approved VPN technology to support the standard directory service ports. Although not required for every network environment, data encryption is likely to be a capability that the selected VPN solution provides. The *Network Infrastructure STIG* and the *Enclave STIG* provide reference information on PPS and VPN use.

- (DS00.1140: CAT II) *If a directory service implementation spans DoD enclave boundaries, the IAM will ensure directory data flowing between directory servers is routed through an encrypting VPN or other network configuration that is compliant with the Network Infrastructure STIG and DoDI 8551.1.*
- (DS05.0130: CAT II) *If a directory synchronization implementation involves the use of LDAP or HTTP across DoD enclave boundaries, the IAO will ensure directory synchronization data is routed through an encrypting VPN or other network configuration that is compliant with the Network Infrastructure STIG and DoDI 8551.1.*

- *(DS05.0140: CAT II) If a directory synchronization implementation involves the use of LDAPS or HTTPS across DoD enclave boundaries, the IAO will ensure directory synchronization data is routed through a configuration that is compliant with the restrictions specified under DoDI 8551.1, such as a DMZ configuration and traffic filtering or a VPN solution compliant with the Network Infrastructure STIG.*

**NOTE:** When directory service implementations employ potentially vulnerable protocols across DoD enclave boundaries, it is especially important to comply with the registration requirements specified under DoDI 8551.1. This helps to ensure DoD network configuration changes do not inadvertently result in a loss of connectivity.

### 3.1.4. Software Integrity

Preserving the integrity of directory data is linked to preserving the integrity of the software that manipulates that data and the directory service environment. Most directory services include both unprivileged user and administrative programs for accessing directory data. While these programs may include a check to verify that the calling user is privileged such as a Windows Administrator or UNIX superuser, some programs may not and others could have flaws that allow such a check to be bypassed. Because of this it is necessary to focus on the access controls for this software as a defense-in-depth measure.

- *(DS00.1150: CAT II) The SA will ensure access to user and administrative directory software libraries (including executable and configuration files) is limited so that:*
  - *Update access is restricted to privileged system accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX) and optionally to a directory service account that is not used to execute the server.*
  - *Read and execute access is restricted to privileged system accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX), authorized application processes, and other IAO-approved users within applicable license agreements.*

The technology-specific appendices of this document list directory service program and configuration files and access control specifications necessary to meet this requirement.

Beyond the software that is bundled with various directory server products, there are additional vendor and third-party synchronization products that manipulate the directory environment and data. Unauthorized access to these programs might allow a user to compromise the confidentiality, integrity, or availability of the directory data or environment.

- *(DS05.0150: CAT II) The SA will ensure access to directory synchronization software libraries (including executable and configuration files) is limited so that:*
  - *Update access is restricted to privileged system accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX) and optionally to a directory service account that is not used to execute the synchronization process.*



- *Read and execute access is restricted to privileged system accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX), authorized application processes, and other IAO-approved users within applicable license agreements.*

In addition to setting appropriate access control permissions, it is a good security practice to perform regular checks to verify that software files have not been modified without authorization and configuration management control. An automated process or tool, referred to as a baseline tool, is used to compare current file characteristics such as date\time, size, and hash to previous values.

- *(DS00.1155: CAT II) The IAO will ensure directory software libraries (including executable files) are checked weekly by an automated baseline process or tool to verify that unauthorized modifications have not been made.*

**NOTE:** This requirement is met for AD by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum*.

- *(DS05.0155: CAT II) The IAO will ensure directory synchronization software libraries (including executable files) used to support production DoD operations are checked weekly by an automated baseline process or tool to verify that unauthorized modifications have not been made.*

The baseline process or tool can be the same as what is used to meet the equivalent requirements for OS files as in the *Windows 2003/XP/2000/Vista Addendum* and the *UNIX STIG*.

Virtually all commercial software requires maintenance to correct vulnerabilities. When a malicious user takes advantage of a vulnerability that is not patched or is never discovered by the vendor, the data and operating environment may be seriously compromised. Because software that is not supported by the vendor could have un-patched or unknown vulnerabilities, it is DoD policy that unsupported software cannot be used.

- *(DS00.1160: CAT I) The IAO will ensure directory service software is removed or upgraded prior to the vendor dropping security patch support.*
- *(DS00.1165: CAT II) The IAO will ensure the site has a formal migration plan for removing or upgrading directory service software prior to the date the vendor drops security patch support.*

**NOTE:** These requirements are met for AD by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum*.

Directory synchronization software is subject to the same potential defects as directory service software. Consequently comparable requirements apply to those products as well.

- *(DS05.0160: CAT I) The IAO will ensure directory synchronization software is removed or upgraded prior to the vendor dropping security patch support.*

- *(DS05.0170: CAT II) The IAO will ensure the site has a formal migration plan for removing or upgrading directory synchronization software prior to the date the vendor drops security patch support.*

Directory software that is used as part of routine operations, whether it is part of the directory service or a directory synchronization solution, is assumed to be providing a function necessary to accomplishing a Component's mission. Unauthorized or improper changes to this software could result in a loss of function. The lack of a documented baseline of the software could be an impediment to recovery in the event of an incident. These risks are addressed by including the directory service and synchronization software in a configuration management process.

- *(DS00.1170: CAT III) The IAO will ensure directory service software that is used to support routine, scheduled DoD operations is included in the baseline inventory maintained by the Configuration Control Board and as part of the Certification & Accreditation documentation, and that a copy of the inventory is stored in a fire-rated container or otherwise not collocated with the original.*

**NOTE:** When the directory service software is tightly integrated into the NOS, as is the case for AD, this requirement is not applicable because the information will be captured as part of the requirements for the host.

- *(DS05.0180: CAT III) The IAO will ensure directory synchronization software that is used to support routine, scheduled DoD operations is included in the baseline inventory maintained by the Configuration Control Board and as part of the Certification & Accreditation documentation, and that a copy of the inventory is stored in a fire-rated container or otherwise not collocated with the original.*

Public domain software that manipulates directory data is readily available through the Internet. The installation or use of this software represents a risk to information systems because the Government does not have access to the original source code to review, extend, or repair it when needed. The software could contain incorrect or intentionally malicious code that would impact the confidentiality, integrity, and availability of the directory service data or environment.

- *(DS05.0190: CAT II) The IAM will ensure binary or machine executable public domain software and other software with limited or no warranty (such as those known as freeware or shareware) is not used to fulfill a directory synchronization function unless the following conditions are met:*
  - *The software is necessary for mission accomplishment and there are no alternative IT solutions available.*
  - *The software is assessed for information assurance impacts and approved for use by the DAA.*

While open source software is not subject to this specific restriction, other requirements are mandated. It is permissible to use open source software as long as it conforms to the same DoD policies that govern COTS and GOTS software. This includes those requirements relating to IA

and IA-enabled components. Specific guidance is found in the DoD Memorandum, *Open Source Software (OSS) in the Department of Defense (DoD)*.

### 3.1.5. Security Service Partitioning

In some implementations, directory servers are used to perform system access and resource authorization functions. Within Windows for example, the AD domain controllers authenticate users and provide data used in access control decisions. In this context these directory servers function as integral parts of the security support structure. When the same host runs other applications, the attack surface of that host is significantly enlarged by the presence of additional programs, data, and application accounts.

Combining a directory server performing access and authorization functions with server applications on the same platform represents a risk that has been determined to be unacceptable for most environments. The following issues contribute to this determination:

- Server applications such as web or database servers typically require a significant increase in the number of installed programs, the number of active processes, and the number of privileged administrative accounts defined. Any of these elements may include vulnerabilities exploited in an attack on the host of the directory server.
- Some applications require the use of application accounts that are members of a privileged group such as Windows Administrators or UNIX superuser. If this is permitted on a directory server, those accounts may have a very broad scope of authority. In Windows, an Administrator on a domain controller becomes effectively an AD Domain Administrator with the potential for significant access to resources throughout the AD forest. If an application account in a privileged group is compromised, the system access and authorization functions performed by the directory service may be compromised as well.
- Some applications require the use of network services on ports that conflict with those used by the directory service. Within AD, the most common examples of this are MS Exchange and ADAM. In these cases, non-standard port numbers are typically used to circumvent this issue. When this is done, traffic to the non-standard ports may not be correctly identified and processed by the host and network intrusion detection systems.
- The use of a common desktop application such as an e-mail client may provide a simple means by which a privileged user unintentionally introduces malicious code on the host.

To address these issues, the use of applications on directory servers that perform system access and resource authorization functions is restricted.

- (DS00.1180: CAT II) *The IAO will ensure directory servers that perform system access and resource authorization functions are not utilized as hosts for applications including database servers, e-mail servers or clients, network address assignment (DHCP) servers, or web servers.*

Note that this restriction does not apply to the following configurations:

- Microsoft Windows Domain Name System (DNS) servers must run on an AD domain controller. As noted in the Domain Name System STIG, the secure implementation of the Microsoft Windows DNS server requires integration with AD.
- Database servers or web servers that are dedicated components required for directory server operation or administration may run on the same host as the directory server, as long as these application servers are not accessed by non-privileged users.

It is recognized that the combination of a directory server with a file server represents less risk than other combinations of applications. In fact, some directory servers require that users be able to access files on the server. Within AD, user and computer accounts must have access to AD Group Policy Template and script files.

However, when the directory server host also acts as a general file server, partitioning directory data from user files (including directory application code) allows more precise access controls to be defined for both. Partitioning also helps to prevent space shortages within user file partitions from affecting the directory service.

- *(DS00.1190: CAT II) If the host of the directory server also functions as a file server, the SA will ensure the directory data files do not reside on the same logical partition as files owned by users.*
- *(DS05.0200: CAT III) The SA will ensure source code for a directory synchronization application does not reside in the same directory as the data input or output to that application.*

### **3.2. Identification and Authentication**

This section describes directory service security requirements based on applicable DoDI 8500.2 IA Controls in the Identification and Authentication subject area. These requirements address items that are used to identify and authenticate a user.

Directory services are commonly repositories of accounts and authentication data used to control access to information systems. In most cases, access to the directory service itself and its authentication data is controlled via the data maintained in it. In some cases, the directory service is also used as a source of data to control network access. Authentication servers that use the Remote Authentication Dial-In User Service (RADIUS) protocol often incorporate or leverage a directory service.

It is noted earlier in this document that version 3 of the LDAP standard was the first version to specify secure directory server authentication controls. Although earlier implementations *could* provide secure authentication, the version 3 implementations *must* provide it. Therefore directory services that enable LDAP access must support LDAP version 3.

- *(DS00.2100: CAT II) If the directory service allows access via LDAP (including LDAPS), the SA will ensure the directory server supports LDAP version 3.*

When a directory provides a repository of authentication data, there must be control over the format and content of that data to make certain that it provides the intended security. For example, a four character password represents weak authentication data that degrades the security of the system on which it is used. It generally enhances security when the control of the authentication data format and content is maintained by the directory server instead of (or in addition to) the applications that update the directory. This ensures uniformity and adherence to standards for updates initiated by any application.

Current policy within DoD is for migration to PKI-based authentication mechanisms. This direction is specified in *DoDI 8520.2, Public Key Infrastructure (PKI) and Public Key (PK) Enabling*, and in other operational documents such as those issued by the JTF-GNO. However, it is acknowledged that some applications do not yet support this direction, but still use passwords as the basis of the authentication data. As long as directory services support those applications, it is necessary to verify that the passwords that reside in the directories meet minimum strength standards.

### **3.2.1. Password-Based Authentication**

The following requirements address the need to enforce stronger authentication data when the directory service uses or supports authentication through passwords.

- *(DS00.2110: CAT II) The IAO will ensure password complexity requirements (length and composition) are enforced for any directory service in which passwords are used or stored, in compliance with the current DoD instructions and JTF-GNO documents and in accordance with current INFOCON status.*
- *(DS00.2115: CAT II) The IAO will ensure password expiration and history requirements are enforced for any directory service in which passwords are used or stored, in compliance with the current DoD instructions and JTF-GNO documents and in accordance with current INFOCON status.*

The following notes apply to the implementation of these requirements:

- These requirements may be met through components of directory server software or application code.
- The requirements must be implemented through all enabled access methods. If application code provides the complexity, expiration, or history controls, this may mean additional controls are needed where network access (e.g., via LDAP on port 389) to the directory can be used to bypass the application.
- If application code is used to implement the password complexity, expiration, or history controls, the following is required:
  - An application security review must be completed for the application.

- It must not be necessary to re-compile the application to change the password complexity, expiration, or history values that are enforced.
  - The expiration period for passwords of directory service process accounts may be up to one year. This includes accounts under which the directory server and replication processes execute.
  - The controls for passwords of directory server process accounts may be administered manually as long as there is documented local policy that addresses compliance, and the accounts controlled under this policy are documented by the IAO. For example, the accounts can be configured without expiration periods as long as a local policy specifies that the passwords are changed within one year, and the affected accounts are listed in IAO documentation.
- *(DS00.2120: CAT I) The SA will ensure all factory set, default, or standard passwords for a directory server application are removed or changed.*
  - *(DS00.2121: CAT III) The SA will ensure all factory set, default, or standard accounts for a directory server application are removed or changed.*

The risk of weak passwords being compromised through password guessing attacks can be mitigated by reacting to multiple failed logon attempts within a defined time period. A requirement to enforce account lockout as specified in the current DoD instructions and JTF-GNO documents is stated in Section 3.3.7, Functional Configuration.

If passwords are stored in a directory database, it is important that access to that data is strictly controlled. Because of the sensitive nature of passwords, control beyond data permissions is required. This additional access control is accomplished through encryption.

- *(DS00.2130: CAT I) The IAO will ensure passwords stored in the directory service database are encrypted.*

**NOTE:** These preceding password requirements are met for AD by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum* and associated documents.

As with other types of access, directory access for synchronization must be authenticated. SAs may choose to automate directory synchronization operations by using scripts that execute under a host account with appropriate access privileges. If the target of the operation is directory data used in identification, authentication, or authorization services, a compromise of the synchronization account could lead to the disclosure, modification, or destruction of sensitive directory data. Even if the data is not used in one of these sensitive services, the compromise of the account could lead to the unauthorized disclosure of directory data.

To protect synchronization application accounts used in scripts from compromise, passwords cannot be embedded in those scripts. If it is necessary to store the password in a secondary file for use by an automated process, the file containing the password must be encrypted.

- *(DS05.0210: CAT I) The IAO will ensure passwords are not embedded in scripts used with directory synchronization solutions and that passwords stored in external files are encrypted.*

As an alternative to files containing account passwords, SAs should investigate whether a script could be converted to run as a Windows service or under a scheduling utility with surrogate capability. This would eliminate the need to store account passwords in external files.

Some directory service technologies employ unique accounts for special functions. The technology-specific appendices of this document provide additional requirements to address password requirements for these accounts.

### **3.2.2. PKI-Based Authentication**

The following requirements address the need to enforce stronger authentication data when the directory service uses or supports authentication through PKI data.

The use of digital certificates can strengthen the identification and authentication process when compared to the use of passwords. However, the use of self-signed certificates and certificates from unauthorized certificate authorities must be avoided. If such certificates are used, it may create a false sense of security, be the source of interoperability issues, or permit circumstances in which unauthorized users gain access to directory data.

- *(DS00.2140: CAT I) The IAO will ensure PKI authentication data stored in a directory database has been issued by the DoD PKI or DoD-approved external PKIs.*
- *(DS05.0220: CAT II) The IAO will ensure PKI certificates used in sessions between directory servers and synchronization clients are issued by the DoD PKI or DoD-approved external PKIs.*

It is noted that certain Components were previously granted a waiver by their CIO to allow the use of non-DoD PKI certificates generated under a Component-wide certificate authority. This action was necessary to support AD domain controllers while the DoD PKI was not able to issue appropriate certificates. This situation has been resolved and Components are now expected to use DoD PKI certificates or have a migration plan in place to implement their use.

### **3.3. Enclave and Computing Environment**

This section describes directory service security requirements based on applicable DoDI 8500.2 IA Controls in the Enclave and Computing Environment subject area. These requirements address eight areas: specific content, architecture and cross-directory authentication, data access control for files, data access control for database objects, data change auditing, group membership and limiting privileges, functional configuration, and data transmission confidentiality and integrity.

### 3.3.1. Specific Content

The content of specific directory data items is generally more an issue of standards than security. But the content of one category of directory service attribute is currently designated for a security function within DoD. This function is known generically in DoDI 8500.2 as the affiliation display.

As with most large enterprises, communication among the DoD Components relies heavily on the use of e-mail. All manner of issues are communicated, discussed, and resolved through messages distributed by e-mail systems. The authors of these messages include military personnel, civilian employees, and contractors; and there may be foreign nationals among those groups. Affiliation refers to an individual's status as a contractor or foreign national.

Certain types of communications are more appropriate or have different authority based on the affiliation of the individual. Consequently it is sometimes significant to know if a particular sender or recipient has a certain affiliation. If an individual is not accurately identified, that person might receive or be excluded from receiving information that is essential to mission accomplishment. For example, affiliation information may be useful in foreign disclosure decisions that are subject to specific public laws and DoD policy.

The schema for each directory may include object attributes to hold e-mail address information. For example, the default AD schema defines attributes for the Contact, inetOrgPerson, and User objects that can be populated with an individual's e-mail addresses. Extensions to the schema may add additional attributes. When these attributes are populated, their values must address the affiliation display requirement.

- *(DS00.0110: CAT II) The IAO will ensure directory attributes used for e-mail addresses and e-mail display names include the abbreviation "ctr" for all contractors and the appropriate country code for all foreign nationals.*

Please refer to the discussion of the Affiliation Display IA control (ECAD) in DoDI 8500.2 for specific guidance on country codes.

DoD policy that specifies the naming convention and acceptable values for some AD User object attributes is detailed in the *Active Directory User Object Attributes Specification*. Although the focus of that document is standards and not security, it impacts the implementation of AD and must be used in conjunction with the requirements in this document.

### 3.3.2. Architecture and Cross-Directory Authentication

When a directory contains identification, authentication, or authorization data, the architecture of the directory service implementation can have an enormous impact to security. Specifically the scope of the directory, in terms of the span of users and resources to which the data applies, must reflect the intended level of trust and control for all the participating users.



In some cases the span of a directory might be deliberately limited in order to isolate groups of users from each other. The need to do this might be as clear as the need to isolate users of DoD contracting data from users of payroll data or as obscure as a local need to isolate development groups within the same branch from each other. This isolation can be one element in achieving need-to-know access control.

It may also be desirable to establish a configuration that supports cross-directory authentication between logically or physically isolated groups. That is, users in one directory may need access to a limited set of resources that are within the scope of another directory. In this case a user could be authenticated through his home directory, and a second authentication using the directory in which the resource resides would not be required. In this way, cross-directory authentication is an element in supporting the concept of single sign-on. Within AD the support for this is known as a trust relationship.

The details of directory implementation architecture and cross-directory authentication configuration depend a great deal on individual product capability and local implementation requirements. This makes it impractical to provide specific requirements that apply generically to directory services. For that reason the requirements and recommendations on these subjects are located in the technology-specific appendices of this document.

### **3.3.3. Data Access Control - Files**

Access to directory data is controlled at two levels. The first of these is the file system object level. This refers to files and directories (the file system component, not the subject of this document) to which access permissions can be applied. The second level of directory data access is the database object level. Depending on the implementation, object permissions that are similar to file permissions can be applied to objects within the database. To preserve the confidentiality, integrity, and availability of directory data, both the file and database object access permissions must be properly configured.

The nature of file system permissions depends on the host OS. On Windows systems this refers to permissions available on NTFS-formatted volumes. On UNIX systems this refers to permission bits and owner\group specifications. In either case access to the directory database files and other files such as log files that are associated with the directory service must be protected. In almost all cases there is no need for users to have access at the file level because they access directory data by communicating with the directory server.

If improper access permissions are defined for directory service files, unauthorized users might be able to directly read, modify, or delete directory data. When this data includes identification, authentication, and authorization data, a compromise could have grave consequences to an entire group of hosts or applications.

- *(DS00.0120: CAT I) The SA will ensure access permissions are configured to restrict access to authorized accounts for the directory database file and any log files, work files, and other associated files that are critical to the directory service.*

The technology-specific appendices of this document list directory service files and appropriate access control specifications necessary to meet this requirement.

Data used in directory synchronization operations has to be considered input to and output from the directory database. In this context, a compromise to the confidentiality and integrity of directory synchronization files could eventually impact directory functions or other operations that use the data. When that data is used for authentication and authorization functions, a serious compromise could result.

Identification of directory synchronization data depends on the product and technology being used. Some examples are:

- The CSVDE and LDIFDE programs can read import files and write export files on AD systems.
- The MIIS product uses a number of files including an MS SQL Server database to hold data being synchronized among directories.

If weak access permissions are defined for these files, unauthorized users might be able to read, modify, or delete data. When this data includes identification, authentication, and authorization data and that data is used to update a directory database, a compromise could have grave consequences to a very large number of applications and hosts.

- *(DS05.0230: CAT I) The SA will ensure access permissions for directory synchronization data files are configured so that:*
  - *Update access is restricted to privileged system accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX) and authorized application processes.*
  - *Read access is restricted to privileged system accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX), authorized application processes, and other IAO-approved users.*

A final issue for data access control for files is related to aggregates of directory data. Aggregates of data outside the directory database are typically the output from or input to synchronization operations. While some aggregates reflect insignificant data, others could effectively disclose sensitive Component force strength or composition data.

If an unauthorized user gains access to a substantial aggregate of directory data contained in synchronization files, that data could be used in an attack or to select valuable targets to attack. While access permissions for synchronization files are required, the added value of the aggregate deserves the additional protection provided by data encryption.

- *(DS05.0240: CAT II) The IAO will ensure directory synchronization files that include a substantial aggregate of the directory data for an entire geographic command are encrypted.*

### 3.3.4. Data Access Control - Directory Database Objects

As mentioned in the previous section, the database object level is the second level at which access to directory data is controlled. However, access control at this level can be very complex because of the need to allow varied access to some directory users while restricting or denying access to other users. Additional complexity arises from the fact that it may be necessary to permit broad read access to a very limited set of directory data. This is so that clients accessing the directory service can examine the root DSE to determine if a desired LDAP control, extended operation, or authentication mechanism is supported.

The nature of database object permissions depends on the directory database implementation. Most directory databases are specialized implementations, designed to support the specific nature of a directory. As a result, the permission implementations vary. Two examples are:

- For AD, directory database objects are assigned access permissions that look much like NTFS file permissions. ACLs for AD database objects are initially created from the default security descriptor for the object type in the AD schema.
- For RHDS, directory database objects can be assigned eight types of access. Permissions are expressed in ACI statements and the permissions are inherited downward from the point in the tree at which they are defined.

If improper access permissions are defined for directory database objects, those objects might be modified, enabling vulnerabilities that lead to immediate unauthorized access or complete disruption of a system. This makes it imperative to properly configure directory database object permissions.

- *(DS00.0130: CAT I) The SA will ensure access permissions are configured to restrict access to authorized accounts for directory database objects.*

The technology-specific appendices of this document list specific directory database objects and appropriate access control specifications necessary to meet this requirement. Generally speaking, the following access control must be enforced:

- Update access to system objects and attributes (typically directory configuration data) must be restricted to administrative personnel and directory process accounts.
- Update access to account (user) objects must be restricted to administrative personnel and directory process accounts.
- Update access to account (user) object attributes that are used for access control must be restricted to administrative personnel and directory process accounts.
- Update access to other account (user) object attributes must be restricted in accordance with the policy for the directory.

One of the critical subjects for directory data access is whether unauthenticated or anonymous connections are allowed. For this discussion the terms unauthenticated and anonymous are assumed to be equivalent and anonymous is used. The following information very briefly discusses the circumstances and guidance for anonymous access.

The nature of directory data is such that read access to it is generally required for a relatively large number of users. Depending on the functions for which a directory service is being used, the data may need to be accessible to users who do not even realize they are accessing it. In the case of Windows, users who are constructing discretionary access rules for files they own enumerate AD user object data in order to build access control entries. Technically these accesses could execute under an anonymous connection to AD. In old versions of Windows OSs, this is how some directory access was accomplished, but this is not true of current, supported versions of Microsoft Windows.

In addition to the concern about anonymous access to typical directory data, there are special considerations for the root DSE entry. Because all directory servers do not support every control, extended operation, and authentication mechanism defined in the LDAP standards, some clients try to read the root DSE entry before attempting to use a specific extension. In executing this query, anonymous access might be attempted. While it may seem harmless to allow this, malicious clients may attempt to take advantage of anonymous access to information in the root DSE to develop attack strategies to compromise the directory service.

Using current directory server and client software, the only circumstance in which anonymous directory access should be allowed is for directories designated as being at the public confidentiality level, as defined in DoDI 8500.2. That is, for directories containing data that anyone is allowed to access. This is expected to be a very small set of cases within DoD. One example might be a directory containing the names and locations of Component recruiting offices.

In recognition that anonymous access to directory data is not an operational necessity for non-public data, a specific requirement to prohibit anonymous access is defined.

- *(DS00.3130: CAT I) The IAO will ensure anonymous access to directory data (outside the root DSE) is not enabled on any directory that contains non-public information.*

Given the content, anonymous access to the root DSE represents considerably less risk compared to other data. Because the access control issues are significantly different, a unique requirement applies to root DSE access.

- *(DS00.3131: CAT III) The IAO will ensure anonymous access to the root DSE of the directory is not enabled on any directory that contains non-public information.*

Some directory service technologies incorporate multiple controls over anonymous access. The technology-specific appendices of this document provide additional requirements to address these controls.

A special class of directory data is the directory schema. As noted in Section 2.2.2, Directory Database, product implementations that are compliant with the LDAPv3 standard store the schema for a directory database as attributes in the database itself. Because of the serious implications of improper or malicious modification to the schema, a special requirement to restrict update access is needed.

- *(DS00.3140: CAT I) The IAO will ensure update access to the directory schema data is restricted to privileged system accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX), dedicated directory management accounts, and authorized application processes.*

A final special issue for directory data access is the use of proxy authorization. This capability was discussed in Section 2.2.1.1, Authentication. Although proxy authorization is an efficient and valuable mechanism for enabling application access to a directory, it can represent a serious vulnerability to the confidentiality or integrity of the directory. If an unauthorized user is granted or gains proxy authorization permission, a large volume of directory data might be improperly disclosed or modified.

A simple review of proxy authorization permissions is not sufficient to determine if the intended resource access policy is being enforced. It is necessary to maintain a baseline of information about valid proxy authorization permission assignment. By comparing the documented baseline to the specific configuration values, improper proxy authorization assignment can be exposed.

- *(DS00.3150: CAT III) The IAM will limit the number of accounts and document those accounts assigned proxy authorization permission.*

### **3.3.5. Data Change Auditing**

In the event of a system compromise, the existence of appropriate audit data can be critical to understanding the extent and significance of the damage. The ability to select the appropriate remedial actions may depend on a review of audit data. It is also important to collect and retain audit data to be able to track and verify the actions of authorized users. In the event of unintended configuration errors, audit data may indicate the source of the error.

If settings are not configured properly to audit changes to objects in the directory database, it may not be possible to determine the source and extent of unauthorized intrusions and unintentional configuration errors.

- *(DS00.0140: CAT II) The SA will ensure auditing is properly configured for the objects in the directory database.*

Because the impact of audit settings depends on activity in the local environment, SAs are advised to monitor the destination for the audit data for their directory database and make adjustments for any increase in the amount of log data.

As actions on specific directory data files and database objects must be audited, the same is true for directory synchronization data. This extends the path of accountability to the data input to and output from the directory database. If directory synchronization operations are not audited, it may not be possible to determine the source and extent of unintentional configuration errors and unauthorized intrusions.

There are four considerations to the implementation of directory synchronization auditing:

- Programs performing synchronization functions must be configured to collect audit data.
  - Programs must be available that allow the audit data to be reviewed.
  - The audit data must be protected from unauthorized access.
  - The audit data must be protected from premature destruction.
- *(DS05.0250: CAT II) The IAO will ensure directory synchronization applications are configured to capture audit-related data.*

For applications running on Windows platforms, the Windows event logging facility is the preferred solution as its use will result in the integration of directory service audit data with other audited events on the system.

Capturing directory database audit data is of little value if that data cannot be reviewed when access to it is required. To meet this need for availability, it is necessary to address review capability, access protection, and retention.

- *(DS00.3170: CAT III) The IAO will ensure tools necessary to review directory database audit data are available.*
- *(DS00.3175: CAT III) The IAO will ensure directory database audit data files are backed up at least weekly onto a different system or media than the system being audited.*
- *(DS00.3180: CAT III) The IAO will ensure backups of directory database audit data files are retained for at least one year.*
- *(DS00.3185: CAT II) The IAO will ensure the permissions for directory database audit data files restrict access to the directory server account, system maintenance accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX), and the local auditors group.*

If the auditing solution for a Windows server-based directory database utilizes the Windows event logging facility and the existing requirements in the *Windows 2003/XP/2000/Vista Addendum* are met, the requirements here for audit data review capability, access protection, and retention are already met.

As is the case with the database audit data, the availability of directory synchronization audit data must also be protected.

- *(DS05.0260: CAT III) The IAO will ensure tools necessary to review directory synchronization audit data are available.*
- *(DS05.0270: CAT III) The IAO will ensure directory synchronization audit data files are backed up at least weekly onto a different system or media than the system being audited.*

- (DS05.0280: CAT III) *The IAO will ensure backups of directory synchronization audit data files are retained for at least one year.*
- (DS05.0290: CAT II) *The IAO will ensure the permissions for directory synchronization audit data files restrict access to the directory synchronization account, system maintenance accounts (such as SYSTEM and Administrators on Windows and superuser on UNIX), and the local auditors group.*

If the auditing solution for a Windows-based directory synchronization solution utilizes the Windows event logging facility and the existing requirements in the *Windows 2003/XP/2000/Vista Addendum* are met, the requirements here for audit data review capability, access protection, and retention are already met.

### 3.3.6. Group Membership and Limiting Privileges

The ability to create groups of user accounts, assign similarly privileged users to them, and use the groups for access authorization is one of the significant benefits offered by a directory service. In many implementations, it is possible to nest groups within groups to extend this benefit. Using groups is an enabler for the security best practice of role-based access control.

Note that the term group is used generically in this discussion to describe a directory object that is a collection of users, other groups, or some combination. Some products use the term security group, role, or something else to describe the same function.

Directory services enable the assignment of special privileges (including “rights” in some implementations) to users and groups. Because of the potential to impact the confidentiality, integrity, and availability of the directory and related applications, it is necessary to control privilege assignment. As an example, restricting sensitive Windows user rights is among the most important tasks in strengthening the security posture of an AD forest and domain. Permission to use many of these rights is the characteristic that distinguishes privileged users or groups. Requirements related to Windows user rights are covered in the *Windows 2003/XP/2000/Vista Addendum* and the *Windows Server 2003 Security Guide*, and elsewhere in this document for directory-specific rights. This section addresses group requirements and strategies for limiting privilege assignment as it applies to groups.

As with the user definitions in a directory, groups are commonly used to control access to the directory itself, as well as for the applications that reference the directory. To secure access control, there are two areas of significant concern for groups:

- Control of the ability to alter group membership
- Control of membership in groups with directory service privileges.

When groups are used for access control, the ability to alter membership in each group must be strictly controlled. Without such control, any resource to which a group has access cannot be considered secure. Since groups are directory database objects, control over group membership is achieved by access controls over the defining objects. By adhering to the requirements in Section 3.3.4, Data Access Control - Directory Database Objects, and the related sections in the technology-specific appendices, the ability to alter group membership is generally controlled through inherited access permissions.

Because special privileges to directory functions and data are commonly assigned at the group level, membership in those privileged groups must be specifically controlled. There are always some groups defined in a directory that have privileged access to that directory. This privileged access may include the ability to alter the configuration of the directory as well as sensitive directory data.

There may be locally-defined privileged groups; but there are always groups defined by the directory service product. The scope of control of these groups varies, but the security of the directory service is directly related to the control of membership in these groups. Ensuring that the membership in privileged groups is controlled requires the maintenance of baseline documentation and periodic reviews to determine that unauthorized users are not members.

- *(DS00.3190: CAT II) The IAM will limit the number of users and document those users assigned to locally-created groups with privileged access to the directory at large.*

Because individual products use different group names, the technology-specific appendices of this document list the product-unique privileged groups and appropriate membership requirements.

The concept of cross-directory authentication is discussed in Section 3.3.2, Architecture and Cross-Directory Authentication. When this capability is considered in conjunction with privileged group membership, special attention is needed. As an example for Windows, the establishment of an AD forest trust makes it possible for a user defined in the trusted forest to be a member of the privileged Domain Admins group in the trusting forest.

While cross-directory authentication may be reasonable for certain data access situations, it is not generally an acceptable practice for administrative privilege assignment. This is because the isolation provided by the separate directories is being effectively eliminated by the cross-directory configuration. If an account in one directory is added to a privileged group in another directory and that account is compromised in the original directory, the compromise could be extended into the scope of the second directory.

There are some isolated instances when cross-directory authentication for a privileged group is acceptable. In AD, the establishment of an account forest and a resource forest within a single organization is such an implementation. A typical case would be an MS Exchange architecture designed to isolate e-mail from other functions of the same organization. Selected Windows users in the account forest could be members of the Administrators group in the Exchange



resource forest as long as both forests are managed by the same organization and comply with the same security policy.

Although there are acceptable cases, the configuration of cross-directory authentication with privileged groups might be indicative of a violation of the separation of duties principle or evidence of an intrusion. Therefore such a configuration must be limited to specific conditions.

- *(DS00.3200: CAT II) The IAO will ensure privileged directory groups do not contain groups or users from another directory unless both directory services are subject to the same security policy and under the control of a single organization.*

It is acknowledged that there are cases when a single person requires administrative privileges in more than one directory and those directories are not subject to the same security policy. This is likely to occur when a perimeter directory is being used to allow access to resources in the perimeter to an outside organization. In this case the better strategy is to assign multiple individual accounts in the two directories. By using an individual account in each directory, the risk is reduced that a compromise in one will spread to another.

An effective way to limit the scope of privileges assigned to accounts and groups is to organize directory objects in ways that allow the use of delegated administration. This can improve directory security because fewer users have privileges on specific collections of objects. For example in Windows, permissions to own or update certain types of AD objects such as OUs can be a delegated administrative privilege.

Delegation strategies and applicable directory objects have to be discussed in specific implementation terms. Refer to the technology-specific appendices of this document for recommendations and requirements.

Another area of concern for limiting privileges is related to the credentials under which the directory server processes execute. In cases where the directory service is tightly integrated with the NOS, the processes may run as part of a system task without a unique account. However, other implementations require an account to be assigned to the directory server and sometimes to other components as well. There are possible security issues created by the groups in which the account is a member and the use of the account for other functions.

There is a tendency to assign membership in the Windows Administrators group or UNIX superuser authority to accounts used for directory service processes or for synchronization operations. This may occur because it is easier than assigning the individual permissions or because the software is not written with the proper approach to security. However, this usually results in granting greater privileges than actually needed and therefore violates the principle of least privilege. If a directory server or synchronization account is assigned greater privileges than required, the processes may be able to read, change, or delete directory data files, directory database objects, or other system objects for which the account was not authorized.

- *(DS00.3210: CAT I) The IAO will ensure accounts used by directory servers or other directory service processes are configured with the least privileges technically feasible. Specifically these accounts will not be members of Windows built-in administrative groups within a domain or assigned UNIX superuser privileges unless no alternative access control permissions can be configured.*

**NOTE:** Some UNIX-based directory server products require superuser privileges in order to have access to the LDAP and LDAPS network ports (389 and 636). Although undesirable, this configuration is currently unavoidable.

- *(DS05.0300: CAT I) The IAO will ensure accounts used by synchronization operations to access directory data are configured with the least privileges technically feasible. Specifically these accounts will not be members of Windows built-in administrative groups within a domain or assigned UNIX superuser privileges unless no alternative access control permissions can be configured.*

Because directory service process accounts and synchronization accounts have a privileged level of access to directory data and possibly system services, the use of those accounts for other functions or as individual user accounts can be a vulnerability on two levels. First it could enable a user to obtain unauthorized access to the directory. Second, it could result in the account being assigned other, unrelated privileges. In both cases, the use of the account can increase the risk of damage if the account is compromised and used for malicious purposes.

- *(DS00.3220: CAT II) The IAO will ensure accounts used by directory servers or other directory service processes are dedicated to that purpose.*
- *(DS05.0310: CAT II) The IAO will ensure accounts used to access directory data in production directory synchronization operations are dedicated to that purpose.*

### **3.3.7. Functional Configuration**

In terms of controls that can be configured for directory servers and directory synchronization programs, there are some items that do not fit easily into the other subsections of the Enclave and Computing Environment controls. This section discusses those items and the associated requirements to enhance security.

The concept of directory replication was discussed in Section 2.2.3, Replication. As noted there, replication is important for providing directory data availability, and can be critical to providing current access control data in environments with multiple directory servers.

Replication scheduling is a directory server configuration item that has a significant impact on replication effectiveness. If a change to access control data is not propagated on a timely basis, users might gain unauthorized access to the directory or systems and data that rely on the directory service.

- *(DS00.3230: CAT II) The IAO will ensure replication is configured to occur at least daily for multi-server environments in which the directory contains identification, authentication, or authorization data.*

Because implementations of replication vary considerably, the technology-specific appendices of this document list specific configuration settings necessary to meet this requirement.

Some directory service implementations have configurable options to control referential integrity in the directory database. When the referential integrity options impact data related to identification, authorization, or authentication, proper configuration is important to ensure current and accurate access control data. For example in RHDS, the failure to enable referential integrity processing could result in outdated group entries that allow access that is no longer authorized.

- *(DS00.3240: CAT II) The SA will ensure directory server or database options that enforce referential integrity for identification, authentication, or authorization data are properly configured.*

Configuration controls that limit successive failed logon attempts are valuable in mitigating the threat posed by password-guessing and other attacks using false authentication data. By locking out accounts for which multiple, invalid authentication attempts are made, it reduces the risk that an account can be compromised.

Depending on the directory service implementation, it may be more efficient to implement account lockout controls through the front-end application that provides access to the directory rather than through control settings defined in the directory server software. Either approach is acceptable as long as the equivalent function is provided for all password-based accounts.

- *(DS00.3250: CAT II) The SA will ensure, for any directory service that supports authentication with passwords, there are controls configured to lock an account after multiple, unsuccessful logon (bind) attempts as specified in the current DoD instructions and JTF-GNO documents, and in accordance with current INFOCON status.*

The following notes apply to the implementation of this requirement:

- The requirement is met for AD by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum* and associated documents.
- If application code is used to implement the lockout controls, the following is required:
  - An application security review must be completed for the application.
  - Replication of the account lockout status within a limited time period must be supported if the account is defined in multiple directory databases.
  - It must not be necessary to re-compile the application to change the invalid logon count that triggers lockout.
- The requirement must be implemented for all accounts, including directory server process accounts. If application code provides the lockout controls, this may mean additional

controls are needed for the process accounts that do not access the directory through the application.

- The requirement must be implemented for all enabled access methods. If application code provides the lockout controls, this may mean additional controls are needed where network access (e.g., via LDAP on port 389) to the directory can be used to bypass the application.

As with other applications, directory servers depend on OS services. In some OS configurations, the services on which the directory server depends are optional. That is, the services may not be running at all times. For example in Windows, services could be configured with a startup type of manual.

The availability of a directory server could be disrupted if required OS services are reconfigured so that they do not always start. Whether this is the result of an attack, error, or a configuration choice, it may be difficult to detect. In order to reduce the probability of a directory server startup error, it is necessary to configure the supporting services to start automatically.

- *(DS00.3260: CAT II) The SA will ensure OS services on which the directory server depends are automatically started.*

Directory server and synchronization solutions generally consist of one or more COTS or GOTS products with predefined and configurable functions. A requirement is stated in Section 3.1.1, Product Design, for a formal security evaluation of those products that maintain identification, authentication, or authorization data. However, in many cases it is possible to alter or supplement the functions that these products perform by adding locally written programs or changing programs supplied with the product. These updates may be enabled through application program interface calls or plug-in configuration statements.

Even though such changes are intended as valuable enhancements, their addition to a previously validated collection of functions could cause unintended, negative consequences. If changes are made to functions that update data used in identification, authentication, or authorization services, the potential exists for the compromise of that IA-related data. In order to ensure the locally written programs or changes are not implemented without adequate review or approval, a configuration management process must be followed.

- *(DS00.3270: CAT III) If a directory server solution is altered by the addition of locally written programs or changes to COTS or GOTS programs, and those programs maintain identification, authentication, or authorization data, the IAO will ensure a formal configuration management process exists and includes approval and review of the implementation of the programs.*
- *(DS05.0320: CAT III) If a directory synchronization solution is altered by the addition of locally written programs or changes to COTS or GOTS programs, and those programs maintain identification, authentication, or authorization data, the IAO will ensure a formal configuration management process exists and includes approval and review of the implementation of the programs.*

### 3.3.8. Data Transmission Confidentiality and Integrity

When data is transmitted over networks, it may be subject to vulnerabilities that allow the confidentiality or integrity of the data to be compromised. Compliance with requirements in the *Enclave STIG* and the *Network Infrastructure STIG* provides protection against many vulnerabilities, but there are actions that can be taken at the directory application level to further strengthen security.

Directory data is transmitted across networks in many instances. This consists primarily of query/update, replication, synchronization, and administrative configuration data that traverse connections using LDAP as well as other protocols. It is essential that the confidentiality and integrity of this data be maintained in order to ensure the integrity and availability of directory data and the applications that depend on it are preserved. In the case of Windows, the integrity and availability of Windows domain controllers, member servers, and clients all depend on the security of the AD data.

Directory replication data from a directory system at the sensitive confidentiality level that is transmitted over wireless connections or over networks that are not subject to DoD controls requires encryption to ensure the confidentiality of the data. If an unauthorized user intercepts the data, sensitive information such as the names and locations of host or application accounts and data resources could be disclosed.

- *(DS00.3280: CAT II) The IAO will ensure replication data is encrypted when it traverses wireless or non-DoD networks.*

**NOTE:** The preceding requirement is met for AD by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum* and associated documents.

Directory replication data from a directory system at the classified confidentiality level may require separate encryption steps if other data on the network is cleared to a lower level or when Sources and Methods Intelligence (SAMI) information is included.

- *(DS00.3281: CAT II) The IAO will ensure replication data is separately encrypted using NSA-approved cryptography when data from a classified directory system traverses a network that is cleared to a lower level than the directory data being transmitted or when SAMI data is being transmitted.*

Directory synchronization data requires confidentiality for the same reasons that directory database data does. Because implementations vary, even those that operate with AD, the following requirements apply to all synchronization solutions.

- *(DS05.0330: CAT II) The IAO will ensure directory synchronization data is encrypted for transport over wireless or non-DoD networks.*

- *(DS05.0331: CAT II) The IAO will ensure directory synchronization data is separately encrypted using NSA-approved cryptography when data from a classified directory system traverses a network that is cleared to a lower level than the directory data being transmitted or when SAMI data is being transmitted.*

Directory configuration data that flows between directory administrators on client machines and directory servers must be protected over all network connections. This reflects the fact that the corruption of this data could seriously compromise the confidentiality, integrity, or availability of the directory server and any applications that depend on the data.

- *(DS00.3290: CAT II) The IAO will ensure sessions between clients used by directory administrators and directory servers are encrypted.*

**NOTES:** The preceding requirement is met for AD administrative tools provided by Microsoft by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum* and associated documents.

For administrative applications without integrated encryption capabilities, this requirement can be met through solutions that provide session-level encryption such as Secure Shell (SSH) or a dedicated Terminal Services configuration.

Substantial aggregates of directory data, even when transported over networks with other security controls, must receive special consideration. Aggregates of data outside the directory database are typically the output from or input to synchronization operations. While some aggregates reflect insignificant data, others could effectively disclose too much Component force strength or composition data.

If an unauthorized user intercepts a substantial aggregate of directory data in transmission, that data could be used in an attack or to select valuable targets to attack. Even with network security controls protecting synchronization data, the added value of the aggregate deserves the additional protection provided by data encryption.

- *(DS05.0340: CAT II) The IAO will ensure when a single directory synchronization operation involves a substantial aggregate of the directory data for an entire geographic command, the data is encrypted for transport over any network.*

The instances in which directory data is transmitted over a network were listed earlier in this section. An important integrity aspect of such transmissions is the verification of the participants at both ends of the network session. If one of the participants is actually an attacker masquerading as an authorized entity, it might be possible for directory data to be disclosed to, or corrupted by, the unauthorized user.

Although network features such as DNS provide a basic level of assurance for originators of a connection that they have contacted the intended server, network address features alone are not sufficient for more sensitive directory server connections. To ensure the integrity of the session,

it is necessary for each participant to authenticate the other. This is referred to as mutual authentication.

The need for mutual authentication is related to the types of logical connections used for transporting directory data. There are two general types:

- A server-to-server connection exists between two directory servers which participate in replication. The initiator of a replication operation could be at either server, depending on whether a push or pull strategy is used by the software.
- A client-to-server connection exists between a client performing query, administrative, or synchronization functions and the target directory server.

In the case of server-to-server (replication) connections, mutual authentication makes certain that the replication supplier is an authorized source directory server and that the replication consumer is the authorized recipient directory server.

- *(DS00.3300: CAT II) The IAO will ensure mutual authentication is performed for directory servers participating in replication.*

If the replication implementation utilizes a specific account, there are considerations for the account and its usage. First, it is necessary to use an account dedicated for replication. By not allowing users or other processes to utilize the account, the principles of least privilege and separation of duties are applied.

- *(DS00.3310: CAT II) For implementations that allow an account to be specified for directory replication, the IAO will ensure the account is dedicated to that function.*

The second consideration for a replication account is related to the authentication process. If a password is used and authentication involves transmission of a clear-text version, then the replication session must be encrypted.

- *(DS00.3320: CAT I) For implementations that allow an account to be specified for directory replication and authentication involves clear-text transmission of the account password, the IAO will ensure replication sessions are encrypted.*

In the case of client-to-server (query\update, synchronization, administrative) connections, mutual authentication makes certain that the directory server is an authorized source or target for directory data, and that the client is an authorized user of the directory. The need for mutual authentication depends on factors including the privileges of the client account, the nature of the operation, and the type and importance of data. Mutual authentication is needed for client-to-server connections used for the following:

- **Administration:** Tasks that involve updates to the configuration and operational controls of the server and updates to identification, authentication, and authorization data
- **Proxy-authorized Updates:** Tasks including updates performed under a proxy identity that has a wide scope of access to directory data

- **Synchronization:** Tasks involving operations to synchronize the target directory with some other directory.
- (DS00.3330: CAT II) *The IAO will ensure mutual authentication is performed for sessions in which directory administration credentials are used.*

**NOTE:** The preceding requirement is met for AD administrative tools provided by Microsoft by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum* and associated documents.

- (DS00.3340: CAT II) *The IAO will ensure mutual authentication is performed for sessions in which updates are done under proxy-based authorization.*
- (DS05.0370: CAT II) *The IAO will ensure mutual authentication is performed for clients and servers participating in directory synchronization.*

The mechanisms used for mutual authentication are not expected to be unique for directory servers. However, the following information is provided to assist in the understanding and selection of potential implementations.

A client or server initiating a session can provide identification that is authenticated by the responding server. Subject to other DoD policy on identification, the authentication data may be in the form of ID and password or PKI certificate.

Authentication of the responding server by the initiator is often less straight forward. There are a couple of common approaches available.

- If the session uses the SSL/TLS protocol (generally through LDAPS), the responding server supplies its PKI certificate as part of the protocol. The initiating client or server validates the certificate to authenticate the responding server.
- If the session uses network or application-level encryption such as that provided by a VPN solution or other proprietary protocol, the hosts or network devices that establish the encrypted session authenticate each other. If used in conjunction with the responding server's authentication of the initiator, such a connection provides an acceptable level of assurance of the identity of the responding server.

The use of an SSL/TLS-based protocol (such as LDAPS) combined with the use of a digital certificate for the identification of the session initiator provides the most desirable implementation of mutual authentication. In this case, both participants are being authenticated through a common mechanism that is typically more robust than ID/password solutions and less complex than proprietary protocol solutions.

Policy on PKI usage is stated in *DoDI 8520.2, Public Key Infrastructure (PKI) and Public Key (PK) Enabling*. In the context of network transport integrity, there are two significant requirements: use of DoD PKI certificates and checking the validity of the certificates. The



requirement to use DoD certificates was stated in Section 3.2.2, PKI-Based Authentication. The validity-checking requirement is addressed here.

If servers or clients do not validate the status of PKI certificates, the unauthorized party in the session may be able to collect client authentication or other sensitive data in the directory. Therefore some method of certificate validation is necessary.

- *(DS00.3350: CAT III) The IAO will ensure directory servers that utilize PKI certificates for authentication incorporate Certificate Revocation List (CRL), Online Certificate Status Protocol (OCSP), or other checking to ensure the validity of certificates.*
- *(DS05.0350: CAT III) The IAO will ensure directory synchronization solutions that utilize PKI certificates for authentication incorporate Certificate Revocation List (CRL), Online Certificate Status Protocol (OCSP), or other checking to ensure the validity of certificates.*

In addition to mutual authentication, some directory service solutions offer other mechanisms to enhance data transmission integrity. When available, data signing can be used at the application level to deter man-in-the-middle attacks that attempt to alter data in transit. In instances where encryption is not required, signing provides greater assurance of data integrity.

- *(DS00.3360: CAT III) The IAO will ensure directory servers are configured to use integrity mechanisms such as data signing to ensure the integrity of transmitted data.*

**NOTE:** The preceding requirement is met for AD by conforming to existing SMB packet signing and LDAP data signing requirements in the *Windows 2003/XP/2000/Vista Addendum* and associated documents.

- *(DS05.0360: CAT III) The IAO will ensure directory synchronization solutions are configured to use integrity mechanisms such as data signing to ensure the integrity of transmitted data.*

The need to terminate inactive connections is a common data transmission security issue. If an inactive network connection is allowed to remain open for an extended period of time, the risk from multiple threats is increased. It could allow an unsecured workstation to be used maliciously. It could allow the injection of malicious traffic during a man-in-the-middle attack. It could enable a successful denial of service attack based on the creation of many idle connections.

Some directory service solutions offer the capability to terminate idle LDAP client connections. Better implementations allow the specification of a default limit with overrides for individual accounts. There may be specific cases in which higher inactivity time limits are appropriate, so documented exceptions are acceptable.

- *(DS00.3370: CAT III) The IAO will ensure directory servers are configured with a default to terminate client connections after 5 minutes of inactivity.*

- *(DS00.3375: CAT III) The IAO will ensure individual accounts that are configured with inactivity timeout limits higher than 5 minutes are documented.*

A final data integrity issue for directory services is related to the use of synchronized time. There is a common reliance within various OS, database, and application server functions on a consistent time base. Directory database replication, Kerberos authentication and authorization functions, and effective audit data generation all depend on accurate time references. If the time on one directory server (such as an AD domain controller) is set to a value significantly different from other directory servers or clients, directory data could be lost or corrupted, clients could be denied logon to a server or access to resources, or audit data could have inaccurate timestamps.

- *(DS00.0150: CAT II) The IAO will ensure a mechanism for synchronizing time is implemented on all directory servers.*

Because the Windows Time service is available on all computers running current versions of Windows, it is logical to use it as the standard tool for all Windows-based directory servers. Network Time Protocol (NTP) daemon software is usually available on computers running UNIX or Linux distributions, so it is the logical choice on those platforms. However, there are environments with mixed collections of operating systems, different tool sets, or network protocol restrictions might indicate a different solution. Because alternative tools or methods could accomplish the same objective, the following guidance is recommended.

- The IAO should ensure a single tool or method is used to synchronize the time on all Windows-based directory servers (including AD domain controllers) within the same forest to a common ultimate time source. The Windows Time service, configured to use the AD domain hierarchy, is a recommended solution. The only exception to this guidance is for the server holding the PDC Emulator Flexible Single-Master Operations (FSMO) role in the forest root domain.
- The IAO should ensure a single tool or method is used to synchronize the time on all UNIX-based directory servers that replicate with each other to a common ultimate time source. NTP daemon software is a recommended solution.

When the Windows Time service is deployed in an AD forest using the domain hierarchy as the time source, the forest root domain PDC Emulator has to be configured to synchronize its time with a trustworthy external source. If this is not done, the time for an entire forest could depend on an imprecise server hardware clock or an unreliable, non-DOD server or network.

- *(DS10.0295: CAT II) If the Windows Time service is used, the IAO will ensure the forest root domain PDC Emulator is configured to synchronize its time with a DOD-authorized time source external to the forest in which it resides.*

Time synchronization clients are commonly capable of automatic switching to alternate sources when a primary source is not available. However, such a switch could indicate malicious activity. Logging of time source switch events enables detection of such activity.

- *(DS00.0151: CAT III) The IAO will ensure time synchronization clients capable of switching time sources create a log entry when a switch is made.*

For the Windows Time service, this level of logging is configurable through the “EventLogFlags” Group Policy setting and is enabled by default in Windows Server 2003. But for Windows 2000 Server it requires the addition of Windows registry entries. The registry entries are documented in Microsoft Knowledge Base article 307937.

It should also be noted that the Windows Time service in Windows Server 2003 supports a very detailed level of event logging. This level of logging is primarily useful for troubleshooting problems, but it may be helpful in some environments to implement on the forest root domain PDC Emulator. Please refer to the documentation in Microsoft Knowledge Base article 816043 and the Windows Time Service Technical Reference on the Microsoft TechNet web site.

### **3.4. Enclave Boundary Defense**

This section describes directory service security requirements based on applicable DoDI 8500.2 IA Controls in the Enclave Boundary Defense subject area. These requirements address items that are related to remote access to directory services and data.

All remote access to DoD information systems, including privileged and unprivileged, requires a restricted access path that includes encryption and strong authentication. This reflects the fact that data flow is being permitted into and out of a protected enclave.

The Terminal Services Windows component is one tool that might be used for remote access. Security-related settings that are required when using Terminal Services are described in Appendix B of the *Windows 2003/XP/2000/Vista Addendum*.

A discussion of privileged user remote access is provided in the *Enclave STIG*. In the context of a directory server or synchronization solution, a privileged user is someone authorized to change the configuration of the product or solution or someone authorized to update identification and authentication data in the directory. If an unauthorized user is able to gain privileged remote access, that user may be able to change security controls to allow directory data to be read, updated, or deleted. To mitigate that vulnerability, stronger access controls are required and session logs must be created and reviewed for each session.

- *(DS00.4100: CAT III) The IAO will ensure remote access to privileged functions of a directory server is secured through a managed access control point such as a Remote Access Server (RAS) and increased session security such as that provided through a VPN.*
- *(DS00.4110: CAT III) The IAO will ensure remote sessions for privileged users of a directory server are logged and the logs are reviewed at least weekly.*

- *(DS05.0380: CAT II) The IAO will ensure remote access to privileged functions of directory synchronization solutions is secured through a managed access control point such as a Remote Access Server (RAS) and increased session security such as that provided through a VPN.*
- *(DS05.0390: CAT II) The IAO will ensure remote sessions for privileged users of directory synchronization solutions are logged and the logs are reviewed at least weekly.*

Remote access by non-privileged users also requires stronger access controls. If an unauthorized user is able to gain non-privileged access, that user may be able to gain read access to sensitive Component force strength or composition data held in the directory or directory synchronization files.

- *(DS00.4120: CAT III) The IAO will ensure remote access to user (non-privileged) functions of a directory server is secured through a managed access control point such as a RAS.*
- *(DS05.0400: CAT III) The IAO will ensure remote access to user (non-privileged) functions of directory synchronization solutions is secured through a managed access control point such as a RAS.*

Remote access to directory servers and directory synchronization solutions implies that a user is accessing the directory system through a network path that includes elements outside the control of the host enclave. If an unauthorized user is able to gain access to a vulnerable network segment, that user may be able to intercept the directory data in transit. To address this threat, encryption is required on all remote access sessions.

- *(DS00.4130: CAT II) The IAO will ensure all remote access to a directory server is encrypted.*
- *(DS05.0410: CAT II) The IAO will ensure all remote access to directory synchronization solutions is encrypted.*

To meet the requirements for secure remote access, an encrypting VPN solution is frequently used. While VPN implementations do provide desirable session protection, they can also be used to conceal malicious traffic. A network-based intrusion detection system (IDS) addresses this threat. A discussion of VPNs and IDS functions can be found in the *Network Infrastructure STIG*. The requirement here enforces the specific need for directory traffic carried by a VPN to be examined for intrusive behavior.

- *(DS00.4140: CAT II) The IAO will ensure VPN traffic for directory data is visible to a network IDS.*

### 3.5. Physical and Environmental

This section describes the directory service security requirements based on applicable DoDI 8500.2 IA Controls in the Physical and Environmental subject area. These requirements address a special need for the physical security of certain servers involved in directory service functions.

Physical access restrictions are necessary for all servers. Such restrictions address common physical threats that are caused in normal office environments where power disruptions, spilled liquids, and cabling disconnections can happen accidentally. These restrictions also address attempts by malicious individuals to disrupt the operation of the server or extract data that may be otherwise protected in transit. Requirements to address physical access to servers are included in the DoD OS (Windows and UNIX) security guidance and are applicable to all directory servers.

Protections for certain types of directory servers are even more significant and require additional review. These servers are critical parts of the infrastructure used for identification and authentication. Because the specific server roles that perform these functions vary by implementation, the technology-specific appendices of this document provide the requirements to address the special physical access requirements.

Some directory synchronization solutions do not require a server OS as their host. This could mean that a synchronization system is not subject to the same physical access restrictions as a machine categorized as a server. However, the functions the synchronization system performs can have a direct, significant impact to directory servers. If an unauthorized user is able to obtain physical access to a machine hosting the directory synchronization solution, that user may be able to compromise the directory server or data just as if they had physical access to the directory server. To address this threat, all hosts of directory synchronization solutions must be subject to the same physical access controls as the related directory server.

- *(DS05.0420: CAT II) The IAO will ensure physical access to host machines used to support normal, scheduled directory synchronization operations is no less restricted than access to the servers on which the directory server software runs.*

### 3.6. Continuity

This section describes directory service security requirements based on applicable DoDI 8500.2 IA Controls in the Continuity subject area. These requirements address three general objectives:

- Directory data, directory synchronization data, and local updates to directory service programs must be backed up according to their unique requirements.
- Information that is required to restore and reconstruct the directory service environment such as directory hierarchy and replication structure must be captured.
- Directory service architecture should be implemented in such a way as to reduce recovery requirements and shorten recovery time.

As indicated in these objectives, restoration or reconstruction of a directory server may require certain information that is not available from a data backup. This information ranges from passwords and software inventories to architecture information needed to rebuild the directory service in the proper sequence. If this information is not available with the backup media when one or more directory servers must be restored or rebuilt, it may result in a permanent loss of directory data including identification, authentication, and authorization data.

NOSs and applications that depend on directory-based identification and authentication data are subject to a loss of availability if the directory server itself becomes unavailable. In the case of Windows, an AD domain cannot function properly without valid and current data from the AD database. Corruption or loss of the data effectively disables a domain controller and can immediately or eventually disable an entire domain and forest. To be able to recover when data in any directory implementation is corrupted or lost, the proper data backup must have been done.

The nature of most directory databases is such that conventional file-level backups do not correctly capture a database that may be in use. As with many database implementations, it is necessary to capture multiple files in a known, coordinated state, to create a backup that is valid for directory restore operations.

- *(DS00.0160: CAT II) The IAO will ensure backup procedures properly capture the directory database and related data daily (preferred) or at least weekly on all directory servers.*

The technology-specific appendices of this document list appropriate backup types necessary to meet this requirement for the unique product database implementations.

Directory synchronization solutions might be used in production operations where database replication cannot provide a function to transport needed data. If an incident causes the corruption or loss of directory synchronization data in those environments, current data may not be available when needed.

- *(DS05.0430: CAT II) The IAO will ensure backup procedures capture the directory synchronization data that is used in production DoD operations.*

Although normal system backups (or other recovery strategies) should capture the COTS programs that make up directory servers and directory synchronization solutions, special attention must be focused on capturing any local programs that might not be collocated with vendor programs.

- *(DS00.6110: CAT III) If a directory server solution is altered by the addition of locally written programs or changes to COTS or GOTS programs, the IAO will ensure backup procedures capture the directory server application code.*
- *(DS05.0440: CAT III) If a directory synchronization solution is altered by the addition of locally written programs or changes to COTS or GOTS programs, the IAO will ensure backup procedures capture the directory synchronization application code.*

Some directory service implementations have unique requirements for their directory restore process. In particular, there may be a requirement to maintain a unique account to use in the process. Most notably, AD requires the use of a dedicated account and a special boot process. The technology-specific appendices of this document provide requirements to address unique accounts used in a restoration process.

In addition to accounts, there can be sequence considerations for restoring directory servers. Particularly in cases where directory hierarchy and database replication configurations exist, the restoration or reconstruction of a directory service must be done in the proper order.

Events caused by natural or man-made circumstances can result in the physical loss of many or all directory servers that support a directory service. If this occurs, reconstruction of the environment is necessary. Beyond the backup media required to perform restores, the availability of some directory implementation information can make reconstruction significantly more efficient. If information about the directory implementation structure is not available, it can be very difficult to restore a directory service quickly and effectively.

- *(DS00.6120: CAT III) The IAO will ensure disaster recovery plans include directory architecture details identifying the hierarchy and replication structure covering all directory systems designated as MAC I or II.*

The technology-specific appendices of this document list directory architecture details that must be recorded to meet this requirement.

Please note that a separate plan for directory service recovery is not required. On the contrary, the best implementation would be an integrated plan for the enclave that includes directory services and other essential applications.

At the time recovery is in progress, it is important to be aware of all the items needed to perform the recovery operation. Directory servers that are tightly integrated with a NOS would not need software beyond the NOS media, but standalone directory server software has to be accounted for. It may be very inconvenient or impossible to assemble a list of software during recovery, so an inventory prepared in advance is critical.

- *(DS00.6130: CAT III) The IAO will ensure disaster recovery plans include identification of directory server software that is used to support production DoD operations on all systems designated as MAC I or II.*

**NOTE:** The preceding requirement is not applicable for AD because of its integration with Windows.

For those environments in which directory synchronization solutions are used in production operations, an inventory of the software required to support those operations is needed.

- *(DS05.0450: CAT III) The IAO will ensure disaster recovery plans include identification of directory synchronization software that is used to support production DoD operations on all systems designated as MAC I or II.*

Incidents that require some limited directory recovery are far more common than disasters. The simple failure of a disk can disable a directory server and require some recovery action. Failure in a network component can disable critical communications needed during OS logon, application access, or resource authorization. It is important to understand that there are directory implementation steps that can be taken to lessen the impact of some incidents.

The basic strategy behind simplifying recovery related to a directory service is the quantity, placement, and functional role of directory servers. If these items are not considered, recovery that might otherwise be largely transparent to users could require outages ranging from hours to days.

One of the easiest recovery strategies is simple redundancy. The presence of multiple directory servers that replicate with each other may allow operations to continue while a single server is recovered. In the case of Windows, multiple domain controllers within each domain can allow local identification, authentication, and authorization functions to continue while a failed domain controller is restored.

- *(DS00.6140: CAT II) The IAO will ensure more than one directory server is operational in each directory service environment that contains servers designated as MAC I or II.*

Proper placement and functional role of directory servers depends on the architecture of the directory service solution and the scope of the deployment. The technology-specific appendices of this document provide recommendations to address this concern for enhancing recovery capability through directory server placement.

### **3.7. Vulnerability and Incident Management**

This section describes the directory service security requirements based on applicable DoDI 8500.2 IA Controls in the Vulnerability and Incident Management subject area. These requirements address the need to be prepared for incidents and to ensure vulnerabilities in directory service products are addressed.

In certain circumstances it may be necessary to alter normal operating configurations in order to address a potential or ongoing computer network attack (CNA). Although the details of a specific attack and the desired defense posture vary, potential configuration changes need to be identified in advance of any CNA incident. The approach to defensive actions is established according to the Information Operations Condition (INFOCON) policy. INFOCON guidance can be found in *United States Strategic Command Directive (SD) 527-1*, *Department of Defense (DoD) Information Operations Condition (INFOCON) System Procedures*, and the *Enclave STIG*.



The concept of cross-directory authentication is discussed in Section 3.3.2, Architecture and Cross-Directory Authentication. The nature of this capability is such that its implementation expands resource access from the scope of users defined in a single directory to users in defined two or more directories. This effectively reduces the overall security level provided by discrete directories. AD manual trusts are a prominent example of a cross-directory authentication capability.

There are times when it may be appropriate to temporarily disable this access. If a CNA incident results in the compromise of one directory, resources within the scope of any directories that allow cross-directory authentication from the compromised directory may be at risk. To address this threat in advance, the procedures to disable specific cross-directory authentication configurations are included in the INFOCON response plan. In the case of AD trusts, the response plan may specify that an external, forest, or realm trust be disabled when a certain INFOCON level is invoked.

- *(DS00.7100: CAT III) Based on the determination of the IAM, the IAO will ensure the local \ applicable incident response plan includes procedures to disable cross-directory authentication configurations as the INFOCON posture increases to higher levels of readiness.*

No software products are immune from the identification and exploit of vulnerabilities. Many vulnerabilities have been linked to deficient programming practices and, although those issues have received a lot of attention, serious problems are still discovered. The proliferation of vulnerability and exploit information on the Internet exacerbates these problems by making the information widely and easily accessible. Unfortunately, mitigating action is often not taken, even when a fix has been identified and made available.

Directory server and synchronization products are not unique in this respect. Although they represent a lesser known target than web server software, the probability that they contain vulnerable code is still significant. If a directory server or synchronization product is found to have a vulnerability and the mitigating patch is not applied, an attacker may be able to exploit the vulnerability to compromise the confidentiality, integrity, or availability of the related directory.

To make certain that vulnerabilities are addressed, a formal commitment to security patch implementation is essential. It is not necessary to have a unique policy for directory server or synchronization products, just a policy that covers them. Manual or automated documentation indicating that patches have been applied provides auditable evidence that mitigating action has been taken.

- *(DS00.7110: CAT II) The IAO will ensure all security related patches to directory server software are applied and that completion is documented for each applicable asset.*

**NOTE:** This requirement is met for AD by conforming to existing requirements in the *Windows 2003/XP/2000/Vista Addendum*.

- *(DS05.0460: CAT II) The IAO will ensure all security related patches to directory synchronization applications are applied and that completion is documented for each applicable asset.*

## **APPENDIX A. RELATED PUBLICATIONS**

### **Government Publications:**

Department of Defense Directive 8500.1, "Information Assurance (IA)," 24 October 2002

Department of Defense Instruction 8500.2, "Information Assurance (IA) Implementation," 6 February 2003

Department of Defense Instruction 8520.2 "Public Key Infrastructure (PKI) and Public Key (PK) Enabling," April 1, 2004

Department of Defense Instruction 8551.1, "Ports, Protocols, and Services Management (PPSM)," 13 August 2004

Department of Defense Memorandum, "Open Source Software (OSS) in the Department of Defense (DoD)," 28 May 2003

Department of Defense, "Active Directory User Object Attributes Specification", Version 1.0, April 2005

Defense Information Systems Agency (DISA), "Database Security Technical Implementation Guide"

Defense Information Systems Agency (DISA), "Domain Name System Security Technical Implementation Guide"

Defense Information Systems Agency (DISA), "Enclave Security Technical Implementation Guide"

Defense Information Systems Agency (DISA), "Network Infrastructure Security Technical Implementation Guide"

Defense Information Systems Agency (DISA), "UNIX Security Technical Implementation Guide"

Defense Information Systems Agency (DISA), "Web Server Security Technical Implementation Guide"

Defense Information Systems Agency (DISA), "Windows 2003/XP/2000/Vista Addendum"

Joint Task Force – Global Network Operations (JTF-GNO), "Concept of Operations for Global Information Grid Enterprise Active Directory", Rev 2, 10 November 2005

National Security Agency (NSA), "Guide to Securing Microsoft Windows 2000 Active Directory," Version 1.0, December 2000

United States Strategic Command Directive (SD) 527-1, "Department of Defense (DoD) Information Operations Condition (INFOCON) System Procedures," 27 January 2006

**Vendor Publications:**

Microsoft Corporation, "Active Directory Application Mode Technical Reference (Draft)", April 2004

Microsoft Corporation, "Active Directory LDAP Compliance", October 2003

Microsoft Corporation, "Best Practice Guide for Securing Active Directory Installations", 2003

Microsoft Corporation, "Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part I", 2003

Microsoft Corporation, "Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part II", 2003

Microsoft Corporation, "Threats and Countermeasures: Security Settings in Windows Server 2003 and Windows XP", 2003

Microsoft Corporation, "Windows Server 2003 Security Guide"

Microsoft Press, "Microsoft Windows Security Resource Kit", March 12, 2003

Red Hat, Inc., "Administrator's Guide, Red Hat Directory Server, Version 7.1", May 2005

Red Hat, Inc., "Configuration, Command, and File Reference, Red Hat Directory Server, Version 7.1", May 2005

Red Hat, Inc., "Deployment Guide, Red Hat Directory Server, Version 7.1", May 2005

Red Hat, Inc., "Gateway Customization Guide, Red Hat Directory Server, Version 7.1", April 2005

Red Hat, Inc., "Red Hat Directory Server Installation Guide, Version 7.1", 2005

**Other Publications:**

Internet Engineering Task Force, "Request For Comments 4513, Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms", June 2006

T. Howes, M. Smith, G. Good, "Understanding and Deploying LDAP Directory Services, Second Edition", Addison-Wesley, 2003

## Web Sites:

### Government:

Information Assurance Support Environment (IASE)	<a href="http://iase.disa.mil/">http://iase.disa.mil/</a> <a href="http://iase.disa.smil.mil/">http://iase.disa.smil.mil/</a>
IASE - Ports and Protocols	<a href="http://iase.disa.mil/ports/index.html">http://iase.disa.mil/ports/index.html</a>
Joint Task Force Global Network Operations (JTF-GNO)	<a href="https://www.jtfgno.mil/">https://www.jtfgno.mil/</a> <a href="https://www.jtfgno.smil.mil/">https://www.jtfgno.smil.mil/</a>
Red Hat Security Solutions License Information	<a href="http://iase.disa.mil/netlic.html">http://iase.disa.mil/netlic.html</a>
Vulnerability Management System (VMS)	<a href="https://vms.disa.mil">https://vms.disa.mil</a>

### Vendor:

Microsoft Windows Server 2003 Active Directory	<a href="http://www.microsoft.com/windowsserver2003/technologies/directory/activedirectory/default.mspx">http://www.microsoft.com/windowsserver2003/technologies/directory/activedirectory/default.mspx</a>
Microsoft Windows Server 2003 Active Directory Application Mode (ADAM)	<a href="http://www.microsoft.com/windowsserver2003/adam/default.mspx">http://www.microsoft.com/windowsserver2003/adam/default.mspx</a>
Microsoft Windows Server 2003 Security Guide	<a href="http://www.microsoft.com/technet/security/prodtech/windowsserver2003/w2003hg/sgch00.mspx">http://www.microsoft.com/technet/security/prodtech/windowsserver2003/w2003hg/sgch00.mspx</a>
Microsoft Best Practice Guide for Securing Active Directory Installations (Windows 2003)	<a href="http://www.microsoft.com/windowsserver2003/techinfo/overview/adsecurity.mspx">http://www.microsoft.com/windowsserver2003/techinfo/overview/adsecurity.mspx</a>
Microsoft Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations (Windows 2000)	<a href="http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/activedirectory/maintain/bpguide/default.mspx">http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/activedirectory/maintain/bpguide/default.mspx</a>
Microsoft Lifecycle Dates	<a href="http://support.microsoft.com/lifecycle/search/">http://support.microsoft.com/lifecycle/search/</a> <a href="http://support.microsoft.com/gp/lifesupsp">http://support.microsoft.com/gp/lifesupsp</a>
Microsoft Security Bulletin Search	<a href="http://www.microsoft.com/technet/security/current.aspx">http://www.microsoft.com/technet/security/current.aspx</a>
Microsoft Support Knowledge Base	<a href="http://support.microsoft.com/search/">http://support.microsoft.com/search/</a>
Microsoft TechNet	<a href="http://technet.microsoft.com">http://technet.microsoft.com</a>

## **Web Sites:**

Red Hat Directory Server Documentation      <http://www.redhat.com/docs/manuals/dir-server/>

Red Hat Knowledgebase      <http://kbase.redhat.com/faq>

## **Other:**

Internet Engineering Task Force (IETF)      <http://www.ietf.org/>

Please note that, for the site links above that span lines, it may be necessary to manually paste the complete links into your web browser.

## **APPENDIX B. LIST OF ACRONYMS**

ACI	Access Control Instruction
ACL	Access Control List
AD	Active Directory
ADAM	Active Directory Application Mode
ADSI	Active Directory Service Interfaces
AIS	Automated Information System
CA	Certification Authority
CIFS	Common Internet File System
CIO	Chief Information Officer
CN	Common Name
CNA	Computer Network Attack
COI	Community of Interest
CONOPS	Concept of Operations
COTS	Commercial-Off-the-Shelf
CRL	Certificate Revocation List
CSV	Comma Separated Value
CSVDE	CSV Data Exchange
DAA	Designated Approving Authority
DADIWG	DoD Active Directory Interoperability Working Group
DC	Domain Controller
DHCP	Dynamic Host Configuration Protocol
DISA	Defense Information Systems Agency
DMZ	Demilitarized Zone
DN	Distinguished Name
DNS	Domain Name System
DoD	Department of Defense
DoDI	DoD Instruction
DSDE	Directory Services Data Exchange
DSA	Directory System Agent
DSE	Directory System Agent (DSA) Specific Entry
DSfW	DSML Services for Windows
DSML	Directory Services Markup Language
DSRM	Directory Services Restore Mode
EAL	Evaluation Assurance Level
FIPS	Federal Information Processing Standards
FRS	File Replication Service
FSMO	Flexible Single-Master Operations
FSO	Field Security Operations
GAL	Global Address List
GC	Global Catalog

GIG	Global Information Grid
GOTS	Government-Off-the-Shelf
GPC	Group Policy Container
GPMC	Group Policy Management Console
GPO	Group Policy Object
GPT	Group Policy Template
GUI	Graphical User Interface
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol over SSL
IA	Information Assurance
IAM	Information Assurance Manager
IANA	Internet Assigned Numbers Authority
IAO	Information Assurance Officer
IASE	Information Assurance Support Environment
IAVM	Information Assurance Vulnerability Management
ID	Identifier
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
IIFP	Identity Integration Feature Pack
IIS	Internet Information Services
ILM	Identity Lifecycle Manager
INFOCON	Information Operations Condition
IP	Internet Protocol
IPSec	Internet Protocol Security
IT	Information Technology
JTF-GNO	Joint Task Force - Global Network Operations
Kbps	Kilobits per second
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LDAPS	Lightweight Directory Access Protocol over SSL
LDIF	LDAP Data Interchange Format
LDIFDE	LDIF Data Exchange
LSA	Local Security Authority
MAC	Mission Assurance Category
MIIS	Microsoft Identity Integration Server
MMC	Microsoft Management Console
MMS	Microsoft Metadirectory Services
MOA	Memorandum of Agreement
MS-DS	Microsoft - Directory Service



NIAP	National Information Assurance Partnership
NIPRNet	Non-secure Internet Protocol Router Network
NIST	National Institute of Standards and Technology
NOS	Network Operating System
NSA	National Security Agency
NT	New Technology
NTFS	NT File System
NTP	Network Time Protocol
OCSP	Online Certificate Status Protocol
OS	Operating System
OSS	Open Source Software
OU	Organizational Unit
PAM	Pluggable Authentication Module
PDC	Primary Domain Controller
PKI	Public Key Infrastructure
PPS	Ports, Protocols, and Services
PPSM	Ports, Protocols, and Services Management
RAS	Remote Access Server
RADIUS	Remote Authentication Dial-In User Service
RBAC	Role-Based Access Control
RFC	Request for Comment
RHDS	Red Hat Directory Server
RID	Relative Identifier
RPC	Remote Procedure Call
RSoP	Resultant Set of Policy
SA	System Administrator
SAM	Security Accounts Manager
SAMI	Sources and Methods Intelligence
SASL	Simple Authentication and Security Layer
SID	Security Identifier
SIE	Server Instance Entry
SMB	Server Message Block
SNTP	Simple Network Time Protocol
SOAP	Simple Object Access Protocol
SSH	Secure Shell
SSL	Secure Sockets Layer
STIG	Security Technical Implementation Guide
TCP	Transmission Control Protocol
TCP/UDP	Transmission Control Protocol / User Datagram Protocol
TLS	Transport Layer Security

UDP	User Datagram Protocol
URL	Uniform Resource Locator
VMS	Vulnerability Management System
VPN	Virtual Private Network
WMI	Windows Management Instrumentation
WQL	WMI Query Language
XML	Extensible Markup Language

## **APPENDIX C. ACTIVE DIRECTORY SPECIFIC ELEMENTS AND REQUIREMENTS**

### **C.1. Introduction**

This appendix addresses the technology-specific background and security requirements for AD. The discussion of security background elements provides high-level technical information about aspects of AD that have security considerations. The security requirements section provides guidance on meeting the general directory service requirements for an AD installation, and additional requirements and recommendations that are unique to AD.

Microsoft's brief definition of AD is a technology "...that enables applications to find, use, and manage directory resources (such as user names, network printers, and permissions) in a distributed computing environment". Although there are capabilities designed for interoperation with other platforms, AD is foremost the directory service for Windows environments.

The importance of AD to an organization is linked inherently to the importance of the Windows servers used by that organization. This is obvious once it is understood that AD is virtually inseparable from any current Windows implementation of more than a few users. It is stated previously, and bears repetition, that AD provides a distributed repository for identification and authentication data. As such AD is critical to enabling and securing shared resources such as files, printers, web sites, and database servers that involve information above the public confidentiality level.

Just as for other organizations, the importance of AD to DoD is inescapable. It is noted in the *Concept of Operations for Global Information Grid Enterprise Active Directory* that AD "...provides localized network directory services, access control, and other services to an expansive array of Component systems (databases, file servers, Community of Interest (COI) websites, etc.) and networks that are connected to the DoD Global Information Grid (GIG)." In recognition of this role and the link between AD and identification and authentication services, this appendix discusses the elements of AD that have security considerations and states requirements to provide a more secure AD environment.

As noted in the body of this document, directory server software must provide some level of compliance with LDAP version 3 to provide sufficient authentication controls. In *Active Directory LDAP Compliance*, Microsoft addresses the older IETF LDAPv3 standards. Windows 2000 Server and Windows Server 2003 both support RFC 2829 and 2830, but are not fully compliant with every provision.

It is also important to note that the directory data that resides in AD is directly and indirectly accessed through other protocols. Windows clients use the Kerberos protocol to access domain controller authentication and authorization services that make determinations based on AD data. The Remote Procedure Call (RPC) protocol is used by domain controllers for AD data replication, by clients for access to Group Policy, and by administrative tools for some interfaces in the Active Directory Service Interfaces (ADSI) technology.

## C.2. Active Directory Security Background

Because of the close relationship, it is difficult to draw lines to clearly separate the elements of Windows from the elements of AD. And while the great majority of AD functions run on Windows servers that are “promoted” to be domain controllers, all the Windows member servers and desktop clients that connect to a Windows domain must run some elements that support AD functions. In concise terms, AD runs as a service on domain controllers and all the non-domain controller computers in a domain are clients of that service.

The *Windows 2003/XP/2000/Vista Addendum*, and by reference Microsoft’s *Windows Server 2003 Security Guide*, provide a great deal of security configuration guidance for Windows servers and clients. These references both contain specifications that address individual domain controllers as well. This appendix to the Directory Services STIG extends that information and provides focused security requirements specifically for AD elements for Windows 2000 Server and Windows Server 2003. This section of this document identifies the AD elements that are subject to the requirements in the next section.

This appendix is not intended as a comprehensive source of information on AD. In fact, an attempt has been made to include only information that is thought to be minimally necessary to identify AD elements. Microsoft and many other authors have produced a lot of documentation and many books that are available for reference. It is also noted that this document is not intended as a tutorial. It is assumed that persons attempting to comply with the stated requirements have a sound understanding of Windows and AD.

Throughout this document references are made to the *Windows 2003/XP/2000/Vista Addendum*. That document and the following documents provided primary input and background to this document:

- Microsoft’s *Windows Server 2003 Security Guide*
- Microsoft’s *Threats and Countermeasures: Security Settings in Windows Server 2003 and Windows XP*
- Microsoft’s *Best Practice Guide for Securing Active Directory Installations* (Windows Server 2003)
- Microsoft’s *Best Practice Guide for Securing Active Directory Installations and Day-to-Day Operations: Part I and Part II* (Windows 2000 Server)
- NSA’s *Guide to Securing Microsoft Windows 2000 Active Directory*.

These and other associated documents are listed in Appendix A, Related Publications.

Based on these statements, the following subsections provide brief descriptions of AD elements for which security considerations exist. Following that, the AD-specific security requirements are described in Section C.3, Technology-Specific Security Requirements. Questions about the AD technology should be researched in the referenced Microsoft documentation and web sites.

### C.2.1. Active Directory Functional Level Considerations

To support a variety of environments that have multiple versions of Windows server software, the concept of domain controller functional levels was implemented in AD. A functional level should be thought of as a set of AD capabilities that all the domain controllers at that level can provide. Please note two differences for Windows Server 2000: the term mode was used rather than domain functional level and the concept of different forest functional levels did not exist.

The following characteristics of functional levels are important:

- The levels are progressive such that the “higher” levels require later OS versions on domain controllers and activate more functions.
- Once raised to a higher level, a domain or forest cannot revert to a lower level.
- There are domain functional levels and forest functional levels. All of the domain controllers in a domain must be at the same domain functional level. All the domain controllers in a forest must be at the same forest functional level.
- The Active Directory Domains and Trusts snap-in to the Microsoft Management Console (MMC) is the primary tool to raise functional levels.
- Windows 2000 mixed is the default domain functional level.
- Windows 2000 is the default forest functional level.

The following table summarizes the available levels and the Windows OS that may be running on the domain controllers in the domains and forests at that level.

Type	Functional Level	Supported DC OS		
		NT	2000	2003
Forest	Windows 2000	Y	Y	Y
	Windows Server 2003 Interim	Y		Y
	Windows Server 2003			Y
Domain	Windows 2000 Mixed	Y	Y	Y
	Windows 2000 Native		Y	Y
	Windows Server 2003 Interim	Y		Y
	Windows Server 2003			Y

**Table C-1. Forest and Domain Functional Levels**

The AD elements that are impacted by the functional level and have a security consideration include:

- The Mixed and Interim levels allow the presence of domain controllers running Windows NT Server in the domain.
- Universal groups are available at the Windows 2000 native domain level and above.
- Group nesting is available at the Windows 2000 native domain level and above.
- The SIDHistory feature is available at the Windows 2000 native domain level and above.
- Forest trusts and the Selective Authentication option for forest trusts are available at the Windows Server 2003 forest level.

Universal groups and group nesting are discussed in Section C.2.3, Group Membership. Forest trusts are discussed in Section C.2.4.3, Manually Defined Trusts.

## **C.2.2. Forest and Domain Architecture**

This section discusses AD forest and domain architecture elements in three general areas:

- Domains, trees, and forests
- Replication, sites, GC servers, and FSMO servers
- Service dependencies and AD data files.

### **C.2.2.1. Domains, Trees, and Forests**

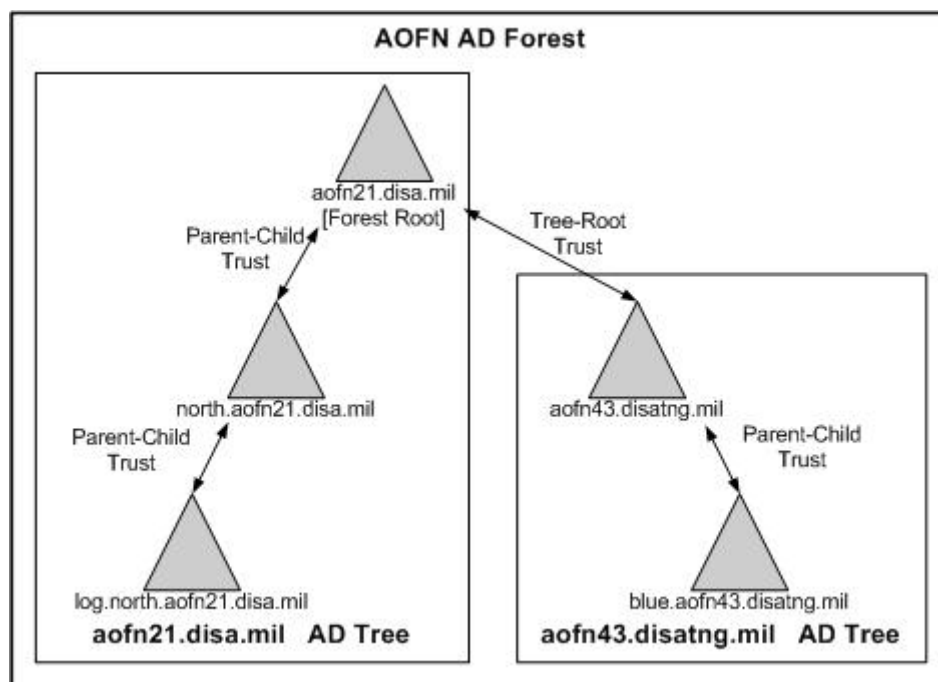
Domains, trees, and forests are terms used to describe hierarchical elements in AD architecture. At their simplest they each represent groups of computers that have different levels of security affiliation. It is important to identify these elements in terms of their security considerations so that related requirements can be understood.

Windows domain terminology was introduced with Windows NT server. Domains are groups of Windows computers that share a common security database for authentication. However, with the introduction of Windows 2000, the meaning of a domain was changed and the implications to security are extremely important. The following distinctions are most significant:

- Windows NT domains represent security boundaries. By default, there are no administrative accounts in one Windows NT domain that have administrative authority in another Windows NT domain.
- Domains in Windows 2000 and above represent administrative boundaries and allow the use of Organizational Units (OUs) to support delegated administration. Some administrative accounts defined in the forest root domain have administrative authority in **all** domains in the same forest. Also, because of the replication function of AD, some changes made by members of the Domain Administrators group are automatically replicated to all the domains in the forest.

Microsoft uses the terms isolation and autonomy to describe domain and forest boundaries. This is a complex subject, but the following greatly simplified distinction is important to security. Isolation refers to a forest boundary; administrative privileges do not cross the forest boundary. Autonomy refers to domain boundaries that allow administrative independence, but some administrative privileges do cross a domain boundary.

The concepts of domains, forests, and trees are best explained in context. The following diagram illustrates a relatively simple forest implementation.



**Figure C-1. Sample AD Forest**

The following characteristics of this forest architecture are important:

- The first domain in the forest, “aofn21.disa.mil”, is known as the forest root domain.
- The “north.aofn21.disa.mil” domain is a child of the aofn21.disa.mil domain. An automatic trust exists between them.
- The “log.north.aofn21.disa.mil” domain is a child of the north.aofn21.disa.mil domain. An automatic trust exists between them.
- The aofn21.disa.mil, north.aofn21.disa.mil, and the log.north.aofn21.disa.mil domains are members of the same AD tree and members of the AOFN forest. They share a common naming context that maps to their names used in the DNS database.
- The “aofn43.disatng.mil” domain is the first domain of a tree that is subordinate to the aofn21.disa.mil forest root domain. An automatic trust exists between them.
- The “blue.aofn43.disatng.mil” domain is a child of the aofn43.disatng.mil domain. An automatic trust exists between them.
- The aofn43.disatng.mil and blue.aofn43.disatng.mil domains are members of the same AD tree and members of the AOFN forest. They share a common naming context that maps to their names used in the DNS database.

There are other AD architectural characteristics that are not depicted. These include:

- An account is defined in only one domain in a forest, but can be used anywhere in the same forest (that resource access permissions permit) because of the automatic trust relationships between domains.

- Security settings that can be managed through the AD Group Policy feature are implemented at the domain level. If the same policy is desired for multiple domains, it must be copied between them.
- A forest root domain is said to be an “empty” domain if it contains only the accounts used to administer that domain and the forest.
- The concept of resource and account domains can be used to partition user definitions from the resources they access. In this model, user accounts are defined in one domain known as the account domain. Resources such as e-mail servers, web servers, database servers, and other application servers are defined in one or more domains known as resource domains.
- When a user in one domain attempts to access resources in another domain in the same forest and Kerberos authentication is used, a domain controller that is trusted by both the user’s computer and resource’s server is used in the authentication process. This is likely to mean that a domain controller in the forest root domain is used in the authorization process.

The AD elements that are impacted by the forest, tree, and domain architecture and have a security consideration include:

- After a user is authenticated in his native domain, he does not have to be authenticated again to access resources in another domain in the forest. This is the effect of the automatic trusts between domains. This provides a kind of single sign-on capability.
- Some security settings including password, account lockout, and Kerberos policy controls, which are implemented through Group Policy, apply to an entire domain. This limits the ability to tailor those settings for individual users or groups.
- In cases where cross-domain resource access is common, placement and security of a forest root domain controller can have a significant impact on the authorization process.
- Implementation of an empty root domain allows stronger security policies to be defined for the sensitive accounts in that domain. It also allows fewer accounts to be defined there, including privileged accounts that might otherwise be needed to support applications. Finally a domain with fewer applications presents a smaller target that might be attacked.
- Where network perimeter protections include a demilitarized zone (DMZ) architecture, the strongest security is obtained by the use of a separate forest for the hosts in that DMZ. This allows fewer network ports to be open because replication traffic is eliminated. It also eliminates the exposure of some information that would otherwise be replicated from the domain controllers on the protected side of the network.

#### **C.2.2.2. Replication, Sites, GC Servers, and FSMO Servers**

Replication, sites, GC servers, and FSMO servers are loosely related subjects, but all have considerations related to architectural choices. From a security perspective, the configuration and placement of these elements can have an important impact to availability and exposure to exploitation.



AD is implemented as a distributed database across all the domain controllers in each AD forest. Some AD data is duplicated on every domain controller in the forest and some AD data is exclusive to all the domain controllers within a single domain. In addition to the Windows directory data held in AD, the settings and information captured for and used by the Group Policy feature are also considered AD data. Other applications, most notably the Microsoft Exchange and Systems Management Server products, utilize AD for storage of their directory-enabled application data.

Replication is the mechanism by which AD data is synchronized among the domain controllers. Although most of the mechanics associated with AD replication require little configuration, there are three considerations related to configuration and security.

- Replication of AD data is handled through the Windows File Replication Service (FRS). This Windows service is mentioned later in this section as a dependency of AD.
- Replication requires network services that transmit data over Internet Protocol (IP) ports. The subject of AD port usage is briefly discussed in Section C.2.6, Ports and Protocols.
- The definition of AD sites impacts replication traffic.

When Windows hosts in an AD forest are distributed across a geographical area that is connected by network links that do not operate at (or close to) Local Area Network (LAN) speeds, it is common to define AD sites. AD site definitions typically mirror physical network boundaries.

Although site definitions are most directly related to network architecture, there are security considerations that must be addressed:

- When sites are defined, it is also necessary to define site links. Site links have properties related to AD replication. The schedule property specifies what hours of the day that a site link can be used. The replication interval property specifies how often, within the schedule period, that domain controllers poll their replication partners to attempt replication. Correctly configured properties ensure replication occurs on a timely basis so that distributed AD security data is kept current.
- An AD site is one level at which Group Policy can be applied. Group Policy is discussed in Section C.2.5, Group Policy.
- When an AD client is logging on to a domain, it attempts to locate a domain controller within the same AD site. A proper site configuration enhances availability and reduces network traffic.

As a directory service AD must support efficient queries concerning objects throughout the span of a forest. The AD GC server was implemented to do this. A GC server is a Windows domain controller that has a full copy of the directory data for the domain in which it resides and a subset of the directory information for every object in the other domains in a forest. The server also has a copy of information that applies forest-wide, including Universal group membership and some AD trust information. From a general replication perspective, the GC can be considered to have a read-only, fractional replication copy of the data in every domain in the AD forest.

GC server location and access have security considerations.

- The GC server functions on domain controllers are accessed through ports 3268 and 3269. The subject of port usage is briefly discussed in Section C.2.6, Ports and Protocols.
- A domain controller accesses the Universal group membership in the GC at the time each user logs on to the domain. In Windows 2000 Server, the inability for a domain controller to contact a GC server could result in denying the user access to the domain. In Windows Server 2003, the Universal group caching function was implemented to reduce the need for synchronous contact.
- The Microsoft Exchange (2000 & later) product and the Outlook e-mail client are dependent on access to a GC server. The unavailability or corruption of the GC server could cause problems in Exchange or Outlook.
- The forest-wide scope of information present on a GC server represents an exploitable source of aggregated data about the forest.

Because AD data is distributed among the domain controllers in a domain, the design of AD includes mechanisms to manage updates from multiple domain controllers. While the design does accommodate updates from multiple sources through the process of multi-master replication, there were some instances in which data integrity required a single-threaded approach. The resolution to this requirement is the implementation of FSMO roles. Please note that FSMO roles are also referred to as operations master roles.

FSMO roles represent specific AD management responsibilities held by assigned domain controllers. Two of the roles apply at the forest level and three at the domain level. AD elements such as AD database schema definitions and certain namespace controls must be managed at the forest level. AD elements such as security identifier (SID) assignment are managed at the domain level. Please note that the subject of FSMO roles is complex and should be thoroughly reviewed in the Microsoft documentation.

The following table summarizes the FSMO roles and lists some of the functions that each performs.

Scope	Role Name	Functions
Forest	Domain Naming Master	Controls the addition or removal of domains Controls the addition or removal of application directory partitions (2003)
	Schema Master	Controls updates to the AD database schema

Scope	Role Name	Functions
Domain	PDC Emulator	Provides a source for time synchronization throughout the domain. At the forest root domain PDC Emulator, provides an authoritative time source for the entire forest Receives preferential password and account lockout updates from other domain controllers and resolves authentication failures due to changed passwords Propagates password changes from down-level clients to other domain controllers Periodically checks and resets ACLs on accounts in certain privileged groups Provides compatibility by acting as Windows NT Primary Domain Controller (PDC) for down-level clients and for Backup Domain Controllers (BDC)
	RID Master	Maintains the Relative Identifier (RID) pool assignments used when a domain controller creates a SID for a new account
	Infrastructure Master	Checks references to objects in other domains in the forest (using a GC server) and maintains the references as changes occur

**Table C-2. Flexible Single-Master Operations Roles**

The AD security considerations for FSMO roles include:

- The availability of the FSMO role holders is directly related to the availability of resources within domains. In some cases, an outage of a FSMO role holder can cause an immediate loss of client access to resources. In other cases, the loss will eventually result in the inability to make changes to AD objects.
- The integrity of the AD database is directly related to the integrity of the Schema Master role holder. While some products such as Microsoft Exchange and Systems Management Server require updates at installation time, there is no need for routine schema updates and any changes should be very carefully considered.
- The Infrastructure Master role is only relevant in forests of more than one domain. Also, placing this role on a domain controller that is also a GC server impedes its function.

### **C.2.2.3. Service Dependencies and AD Data Files**

The final areas of consideration for forest and domain architectural elements are service dependencies and AD data files. Although these are simple, they are significant in terms of a secure AD environment.

In order for AD to function properly there are certain services that must be available and properly configured. Without these services, AD may function improperly or not at all. These services are DNS, a synchronized time source, and Windows FRS.

Access to a secure, properly configured DNS server is a practical prerequisite to AD. Each domain controller dynamically registers DNS service (SRV) records and host (A) records that establish the location of domain controller services including logon and authentication. A Windows server cannot be promoted to be a domain controller without access to a DNS server and a domain controller cannot function without persistent access. Windows clients use DNS to locate domain controllers so the inability of a client to access a DNS server prevents the client from using domain resources. The subject of securing DNS is complex. Please refer to the *DNS STIG* for specific details.

For multiple reasons, it is essential for Windows domain controllers in the same forest to have synchronized time.

- In the AD database objects are stored with timestamps that include date and time of creation and last update. In the event of a conflict detected during replication, the object update timestamp is one factor that may be used to determine which update to retain. As a result the integrity of AD data can depend on an accurate time setting.
- The Kerberos protocol used by domain controllers for authentication and authorization functions requires time on clients and servers to be synchronized. Logon requests from clients with time outside an acceptable tolerance are denied.
- An accurate, synchronized time is critical to auditing functions. Without synchronized time, it may be impossible to correlate the events recorded in the logs of multiple computers. This could make it impossible to correctly evaluate the impact of an intrusion.

Although multiple tools are available to perform time synchronization, the Windows Time service is built into current Windows OSs. It synchronizes server and client computer clocks across a network. The implementation of this service in Windows Server 2003 and Windows XP uses NTP. In Windows 2000 the Simple Network Time Protocol (SNTP) is used. These two Windows implementations use identical network packet formats over User Datagram Protocol (UDP) port 123 and are interoperable for the purposes of time synchronization in AD forests.

For computers within an AD forest, the default configuration of the Windows Time service uses a time source based on the AD domain hierarchy. The Microsoft documentation should be referenced for details, but the following information provides a simplified overview of the default configuration.

- Clients and member servers synchronize their time to the domain controller with which they authenticate.
- Domain controllers within a domain synchronize their time to the PDC Emulator FSMO role holder in their domain.
- For domains outside the forest root domain, the server holding the PDC Emulator role synchronizes its time to the PDC Emulator or any domain controller from its parent domain.
- The server holding the PDC Emulator role in the forest root domain is the authoritative source for time in the forest. It relies on its internal machine clock or has to be configured to synchronize its time to an external time source.

It is apparent from this discussion that the server holding the PDC Emulator role in the forest root domain provides a crucial function for time synchronization. There are two issues that arise as a result. The configuration of a secure time source for that server is critical to ensure an accurate time source for the forest. The designation of a standby operations master, as recommended in Section C.3.5, Continuity, is a very important consideration for continual availability of a time source for the entire forest.

As indicated above, tools other than the Windows Time service are available to perform time synchronization. Although the Windows Time service is highly recommended because it allows a single, consistent method and does not require deployment of a separate program, specific environments may dictate the need for other tools. The most important goal should be an environment in which the ultimate time source is identical for as many servers and clients as possible.

It is noted previously that AD relies on the Windows FRS to replicate data among the domain controllers. AD object data including identity, authentication, and authorization data is moved by way of AD replication. AD Group Policy Template data is also moved through AD replication. AD depends on a functioning replication system to ensure logon and authorization services use current data to make access determinations. It should be noted that it is possible to configure AD to use Simple Mail Transfer Protocol (SMTP) to perform replication of a limited set of database data between AD sites. However, because of the limited scope of this replication and possible propagation delays with network traffic of this type, this is not recommended.

Although the number of AD data files on domain controllers is small, their significance is substantial. The data can be grouped in four general categories of AD data and one category of data related to the FRS service on which AD depends. This data is composed of:

- The primary data store (referred to in this document as the AD database) is named “ntds.dit”.
- The files used by AD for internal transaction logging are named “edb\*.log”, “res1.log”, and “res2.log”.
- Work files used by AD are named “temp.edb” and “edb.chk”.
- Group Policy Template data is stored under the SYSVOL directory.
- FRS data is stored under the Ntfrs directory.

The security considerations for these files are similar to other sensitive files. Integrity and availability are maintained through access control and data backup.

Access control for AD data files is accomplished through the file access permissions available for files on NT File System (NTFS)-formatted volumes. The common SYSTEM and Administrator full control access permissions are assigned.

Because of the way in which Windows accesses AD data, it is not possible to use common file backup methods. AD data must be backed up as part of an operation that backs up the Windows System State data. Among other items, a System State data backup includes the AD database and the Group Policy Template (GPT) data that may be necessary to restore AD.

Conversely, AD data on a domain controller cannot be restored under normal operating conditions. When it is necessary to restore AD data, the domain controller must be booted into a standalone mode called Directory Services Restore Mode (DSRM). This mode can only be entered by supplying the password that is assigned at the time a Windows server is promoted to be a domain controller. This password resides in the Security Accounts Manager (SAM) file on the domain controller.

### **C.2.3. Group Memberships**

The implementation of groups in Windows is the chief mechanism by which Role-Based Access Control (RBAC) may be implemented. By assigning user accounts to groups and referencing the groups in access permissions, users are effectively assigned roles that can be controlled more efficiently and accurately.

A brief note about terminology is important. Windows supports two distinct kinds of groups. The first is the security group. This is a group to which users are assigned for the purpose of access control. The second kind of group is a distribution group. This is a group that can be used by e-mail servers for sending mail to a persistent list of users. In this document, the term group refers to a Windows security group.

Group implementation in Windows is a complex subject that requires substantial documentation to explain well. Readers are directed to the Microsoft documentation and particularly to the *Microsoft Windows Security Resource Kit* for information. The objective of this section is to identify the most significant group issues as they impact AD. The following elements are identified:

- SID assignment and use
- Special privileged groups
- Universal groups
- Group nesting and permission strategies
- OU design
- AD object quotas.

#### **C.2.3.1. SID Assignment and Use**

In Windows each security principal, including users, groups, and computers, is assigned a SID at the time of creation. The SID is a unique value that identifies the security principal within the domain and forest. When a SID is assigned to a new account on a domain controller, the SID includes a RID that makes the account unique within the forest. The RID comes from a pool of values that are assigned to the domain controller by the RID master FSMO server. Thus an account SID can be identified as coming from a specific domain.

A SID is never reused. If an account is moved from one domain to another, a new SID is assigned to the account. Starting at the Windows 2000 native domain functional level, the SIDHistory feature allows old SIDs to be retained along with the new SID that is assigned when the account is moved.

SIDs play a primary role in resource authorization. When a user logs on to his domain, the SID for his account, the SIDs for all groups of which he is a member, and any SIDHistory values are extracted from the AD database. When the user attempts to access a resource, this list of SIDs is compared to the ACL for the resource to determine what access is permitted.

This discussion of SID assignment and resource authorization is basic to understanding a potential vulnerability that involves resource authorization through an AD trust. The discussion of this issue is in Section C.2.4.3, Manually Defined Trusts.

### **C.2.3.2. Special Privileged Groups**

There are several pre-defined Windows groups that can be categorized as specially privileged. This is primarily true because the ACLs of many Windows resources are defined by default with access permitted to those groups. It is also true because some programs examine group membership to determine if the user is allowed to execute sensitive functions in the program.

The following briefly identifies the groups and security privilege implications for AD:

- Domain Admins and Enterprise Admins - Members of these groups have permissions to all AD objects at the domain and forest level respectively.
- Schema Admins - Members of this group have permission to modify the AD schema for a forest. This allows the addition, deletion, and modification of AD object definitions and their attributes. Among the object attributes are the default security descriptors that represent the default access permissions that are assigned to objects created from the schema definition.
- Group Policy Creator Owners - Members of this group have permission to add, delete, or modify Group Policy Objects (GPOs).
- Pre-Windows 2000 Compatible Access - Membership in this group allows users read access to many AD objects. The primary security issue associated with this group is anonymous access to AD data. This occurs when the group membership includes anonymous users or other groups that include anonymous users.
- Incoming Forest Trust Builders - In the forest root domain of Windows Server 2003 domain controllers, members of this group are allowed to create incoming, one-way forest trusts.

The impact of group nesting has to be considered for privileged groups. For example, the Domain Admins group is a member of the local Administrators group on all computers in the domain. This means that any user in the Domain Admins group has privileged access to every computer in the domain. Group nesting is discussed later in this section.

There are mechanisms available within AD to strengthen the security associated with privileged group membership. These include the AdminSDHolder object and the Restricted Groups Group Policy setting.

AdminSDHolder is an AD object that acts as a template for the security descriptor attributes of certain privileged accounts. Every hour the domain controller holding the PDC Emulator FSMO role compares the ACL on the AdminSDHolder object to the ACLs of accounts in certain privileged groups. This includes the Domain Admins, Enterprise Admins, and Schema Admins groups among others. If the ACL on an account differs, the ACL of that account object is overwritten using the ACL of the AdminSDHolder object. This ensures that the permissions on the account objects have not been compromised.

The Restricted Groups security setting in a Group Policy can be used to control the membership of a group. The Restricted Groups setting can be configured with two properties for each group. The “Members” property specifies the accounts that are to be members of the group. An empty Members property specifies that the group will have no members. The “Members of” property specifies which group(s) the subject group should be a member in. When the Restricted Groups setting is defined in a GPO, the group memberships are checked each time Group Policy is refreshed. As an example, if a GPO with a Restricted Groups setting is configured for the Schema Admins group and the Members property is empty, accounts that are inadvertently or maliciously added to the Schema Admins group are removed the next time Group Policy is refreshed.

### **C.2.3.3. Universal Groups**

A “universal” group is not a specific group, but rather one of the types of Windows groups. It can have members (accounts) from any domain in a forest and so can be useful when resource access permissions need to be applied to users in many domains. (Conversely there is no logical need for universal groups in a single-domain forest.) There are several considerations for universal groups for AD security:

- Universal groups are available starting at the Windows 2000 native domain functional level.
- Because they are forest-wide objects, universal groups are kept in the GC and changes are replicated forest-wide. This can have network traffic and propagation delay implications if membership changes are made frequently.
- Each time a user logs on to the domain, the domain controller checks the GC server to retrieve the SIDs of the universal groups that have the account as a member. Unless the Windows Server 2003 universal group caching function is available, unavailability of the GC server prevents logon.
- The use of SID filtering can cause the SID of a universal group to be discarded during resource authorization across a trust. This occurs when the universal group is created in a domain different from the user’s domain. SID filtering is discussed in Section C.2.4.3, Manually Defined Trusts.

### **C.2.3.4. Group Nesting and Permission Strategies**

The concept of group membership for accounts has been discussed, but some brief words about group nesting and permissions strategy is important. Group nesting affects, and is affected by, AD configuration.



Group nesting refers to the ability to embed one group within another. It is available starting at the Windows 2000 native domain functional level. This capability makes it substantially easier and more efficient to construct access permissions. This enhanced ease and efficiency can result in more accurate and therefore more secure permissions being assigned.

Microsoft has suggested some strategies that utilize group nesting when applying permissions. The use of these strategies provides an organized way to implement RBAC. There are three basic strategies; they are referenced by letter sequences:

- A-G-DL-P - In this strategy, accounts are added to global groups, global groups are added to domain local groups, and resource permissions are granted to domain local groups.
- A-G-G-DL-P - This strategy adds nesting of global groups where that helps to reduce the overall number of groups. Accounts are added to global groups, global groups are nested in global groups, global groups are added to domain local groups, and resource permissions are granted to domain local groups.
- A-G-(G-)U-DL-P - This strategy adds universal groups or replaces one level of global groups with universal groups. Accounts are added to global groups, global groups are nested in global groups, global groups are nested in universal groups, universal groups are added to domain local groups, and resource permissions are granted to domain local groups.

Two common characteristics are consistent throughout all of these strategies. Accounts are always added to global groups and permissions are only granted to domain local groups.

#### **C.2.3.5. OU Design**

OUs are mentioned here primarily to ensure the distinction between an OU and a group is noted. An OU is a collection of users or computers that is defined in AD for ease of management. An OU does not have a SID attribute. Groups are assigned SIDs that can be referenced in ACLs that are used to control resource access. OUs are used with GPOs to apply security settings and other configuration settings.

It should also be noted that administration of OUs can be delegated to non-Administrator accounts. This is recommended because it helps to implement the security principles of separation of duties and least privilege.

#### **C.2.3.6. AD Object Quotas**

The final AD element to be discussed in the context of group membership is AD object quotas. This subject is only loosely connected to groups in that two built-in groups are automatically exempt from it.

The AD object quota was introduced in Windows Server 2003. It is a means by which a limit can be imposed on the number of objects that a single account can own in a given AD database partition. This control provides a defense against inadvertent or deliberate attempts to exhaust the

capacity of an AD database. This could occur when a member of the Group Policy Creator Owners group creates a scripted loop that adds GPOs.

The security considerations for AD object quotas include:

- Separate quotas are applied to each AD database partition, except for the schema partition which is exempt from quotas.
- Deleted objects, known also as tombstones, are factored into the count of owned objects. It is possible to alter the factor so that tombstones count for less than one object.
- Quotas can be defined as the default applied to all accounts or for individual accounts.
- When not defined, the default quota is no limit.
- Members of the Enterprise Admins and Domain Admins groups are exempt from quotas.

#### **C.2.4. Trust Relationships**

AD trust relationships are an inherent part of domain and forest architecture. Trusts are the mechanism for allowing a user to authenticate to one domain and access resources in another domain without authenticating again. Trusts automatically exist between domains in the same forest, but can be configured between domains in different forests, between two forests (with Windows Server 2003), and between a domain in a forest and a Windows NT domain or UNIX Kerberos realm outside the forest. When a trust is configured between entities that are not within the same forest, this is an instance of cross-directory authentication as discussed in Section 3.3.2, Architecture and Cross-Directory Authentication.

Every time a domain is defined, the trust configuration within the forest is automatically updated. Identifying trust types and their configuration options is essential for understanding the impact on resource access control. This section first describes certain trust properties and associated options. The following subsections describe the trusts that are created by default (automatically) or by manual administrator action.

##### **C.2.4.1. Trust Properties and Terms**

This section describes some common properties of trusts, terminology used with trust descriptions, and a figure illustrating how trusts might be defined. The following subsections describe the trusts that are automatically and manually defined in AD environments.

There are two properties that significantly impact the effect of every trust. These are transitivity and direction. For some trust types the direction property is fixed; other types allow configuration.

The property of transitivity refers to a logical relationship. It is explained generically by considering the relationship between three elements. If there is a relationship between A and B, and the same relationship between B and C, then transitivity exists if the same relationship exists between A and C. When this is applied to AD trusts, it would mean that: domain A trusts domain B, domain B trusts domain C, and for a transitive trust domain A trusts domain C.

The property of direction refers to one-way or two-way trust flow. It is possible to think of flow as the direction in which user credentials (SIDs) travel for resource authorization. If a user in domain A attempts to access a resource in domain B, and a one-way trust (at least) exists between A and B, the user's credentials travel from A to B. When users from each domain can access resources in the other, a two-way trust exists.

There are two additional sets of terminology that are used in trust descriptions:

- Trusted and trusting - A trusted domain is the domain where the user is authenticated and his SIDs are extracted from AD. A trusting domain is the domain that accepts the user's SIDs as passed from the trusted domain.
- Incoming and outgoing - An incoming trust specifies that the local domain is trusted by another domain. An outgoing trust specifies that another domain is trusted by the local domain.

A final concept that is important for the trust discussion is the trust path. A trust path refers to the number of domain controllers that have to be contacted in order to validate that a trust relationship exists that allows a specific access to occur. This concept becomes significant in a complex forest with an extensive tree structure. In that case, the trust path between a domain "far down" in one tree to a domain far down in another tree can require traversal of several domains in each tree. In Windows Server 2003 domains, a client cannot traverse more than 10 trusts to access a resource.

Trust concepts are best explained in context. The following diagram illustrates where different trust types might be implemented and is useful for the following discussions.

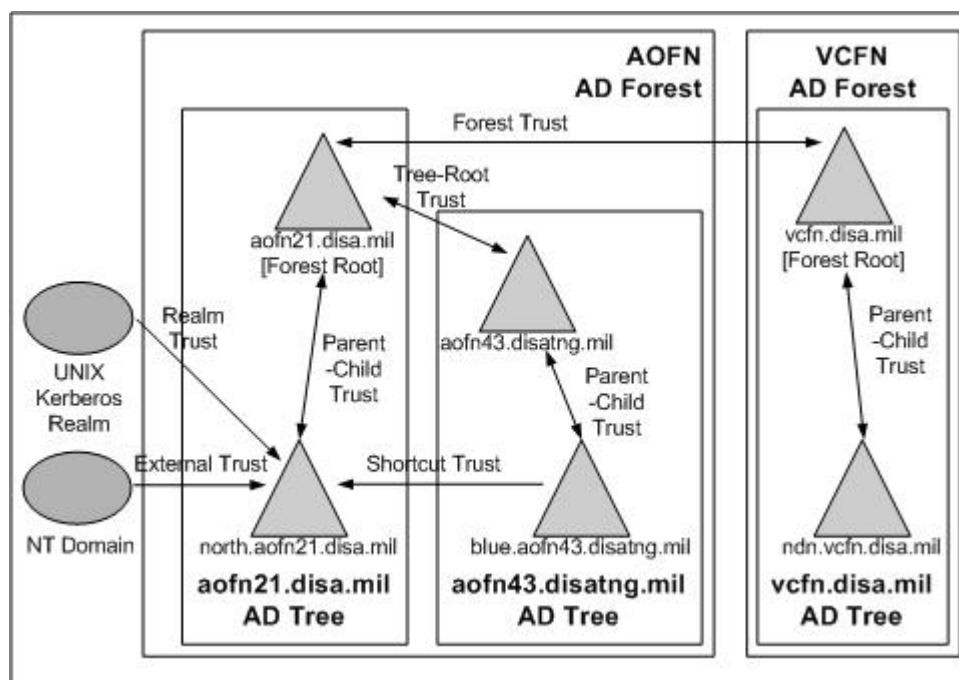


Figure C-2. Sample Trusts

#### C.2.4.2. Automatically Defined Trusts

It is noted previously that adding a domain to a forest automatically alters the trust configuration in the forest. Specifically, as a domain is added to an existing forest, a trust is automatically added between it and the domain to which it is logically connected. This logical connection is in one of the following forms:

- A parent-child trust relationship exists between two domains within the same tree.
- A tree-root trust relationship exists between the forest root domain and the first domain in a different tree in the forest.

Trusts between domains within a forest are identified by relationship type. They are created with the following properties:

Trust Type	Transitivity	Direction
Parent-child	Transitive	Two-way
Tree-root	Transitive	Two-way

**Table C-3. Automatic Trust Types**

These properties reflect the fact that forest design is intended to enable resource access across domains. That is, the two-way, transitive nature of these trusts allows a user in any of the domains to access resources in the other domains to which that user has the appropriate access permissions. This is a primary reason why the way in which forest and domain architecture is implemented has such significant security considerations.

#### C.2.4.3. Manually Defined Trusts

There are situations in which trusts are needed beyond those created automatically within a forest. In most cases this reflects a desire to extend trust beyond the forest boundary, but there is one implementation that is used within a forest for performance reasons.

The following trust types are manually created as needed:

- External - An external trust can be created between two domains in different forests or between an AD domain and a Windows NT domain.
- Forest - A forest trust can be created between the forest root domains of two forests. However, forest trusts can only be created between forests that are operating at the Windows Server 2003 forest functional level. This requires that all the domain controllers in both forests are running Windows Server 2003.
- Shortcut - A shortcut trust can be defined between two domains in the same forest. A shortcut trust is used where the trust path is long or network connections between domain controllers in the trust path cannot efficiently support the authorization traffic.
- Realm - A realm trust can be created between a domain and a non-Windows system such as a system hosting a UNIX or Linux OS with Kerberos version 5.

These trusts are created with the following properties:

<b>Trust Type</b>	<b>Transitivity</b>	<b>Direction</b>	<b>Authentication</b>	<b>SID Filtering</b>
External	Non-transitive	One-way	Domain-wide or Selective	Enabled or Disabled
Forest	Transitive (within forests)	One-way or two-way	Forest-wide or Selective	Enabled or Disabled
Shortcut	Partial	One-way or two-way	N/A	N/A
Realm	Either (With supporting partner)	One-way	N/A	N/A

**Table C-4. Manual Trust Types**

The Authentication property became available in Windows Server 2003. It supports more granular control over trusts. It is implemented through an additional property for external and forest trusts and a new permission for AD computer objects:

- The trust Authentication property defaults to a value that allows all users in the trusted forest or domain to be authenticated via the trust. The other option is Selective authentication.
- The “Allowed to Authenticate” permission for AD computer objects is not set by default. It can be set to Allow or Deny for specific users or groups.

When an external or forest trust has the Selective authentication option set and a user from a trusted forest or domain attempts to access an object, the Allowed to Authenticate permission on the computer object in the trusting domain is checked. If the permission is set to “Allow”, access is permitted.

The SID Filtering property was added in a security patch to Windows 2000 Server and was built into Windows Server 2003. It was designed to prevent an elevation of privilege attack that could result from the way authorization data from trusted forests is used. Information on this vulnerability is documented in Microsoft security bulletin MS02-001.

In Section C.2.3, Group Membership, the concept of account SIDs was discussed. As noted there, when a user logs on, the SID for his account, the groups of which he is a member, and any SIDHistory attribute values are extracted from the AD database. When a user attempts to access a resource through an external or forest trust, this collection of SID authentication data is passed to the domain where the resource exists.

A vulnerability existed because trusting domains did not verify that the SIDs in the authentication data all came from the trusted domain. If the authentication data is compromised by the addition of SIDs known to have access permissions to the resource, then an unauthorized user may gain access that he would not otherwise have.

When the SID Filtering property is enabled, the trusting domain removes any SID in the authentication data that was not generated in the trusted domain. This ensures that only accounts created in the trusted domain are eligible for access to resources in the trusting domain.

Some notes on the use of SID Filtering are important:

- The removal of SIDHistory values from authentication data could cause an unintended denial of access. SIDHistory SIDs are removed because they do not appear to be from the trusted domain.
- SID Filtering can cause the SIDs for universal groups to be removed from authentication data. When the universal group was not created in the same domain as the user, its SID does not appear to be from the trusted domain.
- SID Filtering must only be enabled for trusts that span forest boundaries. An attempt to enable SID Filtering on domains within a forest will cause trust and replication failures.
- The term “quarantine” is sometimes used to describe SID Filtering and is an operand of the Domain Manager (netdom.exe) command line program used to configure SID Filtering.

The following AD security considerations apply to manual trust definitions:

- Establishing a trust relationship outside a forest makes one or both participants dependent on the security practices of the other. In simple terms for a forest trust, the trusting forest is relying on the identification and authentication practices in the trusted forest. Therefore it is important that all the users in both forests who are members of the highly privileged groups (Domain Admins, Enterprise Admins, and Schema Admins) are considered to be highly trusted individuals by both organizations involved in the trust.
- Trusts may be configured through multiple tools. Two common tools are the AD Domains and Trusts MMC snap-in and the Domain Manager (netdom.exe) command line program.
- Whenever possible, one-way trusts should be chosen over two-way. In cases such as perimeter configurations with a separate forest, a one-way trust alone can be used to allow accounts from the protected-side forest to be trusted in the perimeter forest while preventing accounts in the perimeter forest from being trusted in the protected-side forest.

Because of the possible impact on Windows OS security, a brief note on trust configurations and some specific Group Policy settings is necessary. Some Windows NT components allowed and used anonymous access to retrieve Windows data for configuration and communication tasks. The operations involved in the establishment of a trust relationship are one example. The following table describes the registry entries and associated Group Policy settings that are related to this anonymous access.

OS	Registry Entry	Group Policy
Windows 2000 Server	RestrictAnonymous	Additional restrictions for anonymous connections

OS	Registry Entry	Group Policy
Windows Server 2003	RestrictAnonymous	Network access: Do not allow anonymous enumeration of SAM accounts
	RestrictAnonymousSam	Network access: Do not allow anonymous enumeration of SAM accounts and shares

**Table C-5. Anonymous Access Settings**

Secure values for these entries are required by the *Windows 2003/XP/2000/Vista Addendum*. However, it is acknowledged that using the most secure values causes trusts with Windows NT 4.0 domains to fail. For this reason, it is permissible to use the less secure values in those environments. However, once all the domains involved in trusts are at Windows 2000 and above, there is no longer a requirement for these values as far as AD trust support is concerned.

Please note that the reduced security OS settings for anonymous access may be required for other applications in an environment. It is important that those cases are properly documented so that this need for reduced security settings can eventually be resolved.

### **C.2.5. Group Policy**

Group Policy is a Windows feature implemented through AD that has several functions. Very generally speaking, it provides a method to define Windows configurations and enforce those configurations for the computers and users within a Windows domain. Although explicit security settings are one prominent aspect of Group Policy, there are many others.

Group Policy implementation and use are very complex subjects. Microsoft has stated that the Group Policy implementation in Windows Server 2003 has almost 1,000 configurable settings. This statement reflects the fact that a few brief notes cannot adequately explain such a complex technology. The information here is intended only to identify the most significant elements of this technology as it relates to AD security. Readers are strongly encouraged to review the Microsoft documentation including the *Microsoft Windows Security Resource Kit* for information.

It is noted that the discussion here is addressed primarily at how Group Policy is applied through Windows domain controllers. The interaction with locally defined Group Policy is mentioned, but it is not the focus of this information.

#### **C.2.5.1. Group Policy Components**

Group Policy is stored logically in GPO objects in AD. From a physical viewpoint most policies have two data components. The Group Policy Container (GPC) is stored in the AD database. The Group Policy Template (GPT) is stored in the SYSVOL folder on Windows domain controllers.

The physical GPO components impact replication, backup, and restore. AD services manage the replication of GPO data in the AD database. The Windows FRS replicates the GPTs in the SYSVOL folder among all domain controllers in a domain. Synchronization between the AD database and SYSVOL contents is a consideration for backup and restore operations.

### **C.2.5.2. Default GPOs**

When a domain is created (when a Windows server is promoted to be the first domain controller), two GPOs are automatically built. As each subsequent domain controller is created, these GPOs are replicated to the new domain controller. The Default Domain Policy GPO is the default policy linked to the domain; it is applied to users and computers throughout the domain. The Default Domain Controllers Policy GPO is the default policy linked to the Domain Controllers OU; it is applied to all domain controllers in the domain.

### **C.2.5.3. Application of GPOs to Objects**

Implementing GPOs can be complex. In the simplest terms, there are two basic parts to the process: definition and linking. Definition is simply the creation of a GPO and setting its properties. Linking refers to the process of specifying the AD object to which a specific GPO is to be applied. A GPO has no effect until it is linked to an AD object.

GPOs can be linked at three levels of objects:

- GPOs linked to an AD site are applied to all users and computers at that site.
- GPOs linked to the domain are applied to all users and computers in that domain.
- GPOs linked to an OU are applied to all users and computers in that OU.

Two of these levels always apply to an individual user or computer. The user or computer always belongs to an AD site (even if it is just the default site). The user or computer always belongs to an AD domain. Users and computers should belong to an OU, and most GPOs are linked to OUs.

A single user or computer may be subject to multiple GPOs. In addition to this, it is important to understand that GPOs are applied in a specific sequence. This sequence is sometimes referred to as “LSDOU”:

- Local - GPOs defined on the local computer
- Site - GPOs linked to the applicable AD site
- Domain - GPOs linked to the applicable domain
- OU - GPOs linked to all applicable OUs.

Within each of these categories, there may be multiple GPOs. The settings in the GPOs are aggregated to arrive at the final settings for the user or computer. The sequence in which they are applied within each category is determined by the link order that is configured. Also the concept of inheritance applies. Inheritance refers to the fact that policies applied to parent containers are also applied to the child containers. In practice this means that a policy applied to an OU also applies to all the OUs defined within that OU.

There are three GPO options that can be used to change normal GPO application behavior.

- Enforcement (No override) - This option specifies that a GPO takes precedence over any GPOs linked to child containers.



- Block inheritance - This option specifies that objects in child containers do not inherit GPOs from parent containers.
- Loopback processing - This option specifies that the computer GPOs should be applied when any user logs on to the computer. Loopback processing operates in merge mode or replace mode. In merge mode the user parts of the computer's GPO are applied along with the user's OU GPO. In replace mode the computer's GPO is applied instead of the user's GPO.

Filtering can also impact GPO application. Security filtering occurs when a GPO is linked to a computer or user, but the computer or user account does not have Read and Apply permissions for the GPO. An account must have these permissions in order for a GPO to be applied. In Windows Server 2003, Windows Management Interface (WMI) filtering was added. WMI filtering allows the characteristics of a computer to be checked to determine if a GPO should be applied. Criteria are specified through WMI Query Language (WQL) statements.

A final consideration for GPO application is the designation of a slow link. Because GPO application occurs each time a computer connects to a domain, when the user logs on to the domain, and periodically thereafter, the speed of the connection from the client to the domain controller could substantially impact performance. To account for this, the speed of the link from the domain controller to the client is evaluated. If that speed is below a configured threshold (default 500 kilobits per second (kbps)), the link is designated as a slow link and some GPOs are not applied. Those that are not applied include scripts, folder redirection, disk quotas, and application deployment. This behavior can be overridden through a Group Policy setting.

The following AD security considerations apply to the GPOs:

- Linking a GPO to an AD site is restricted to members of the Enterprise Admins group.
- Linking a GPO to a domain is restricted to members of the Domain Admins group.
- Linking a GPO to an OU can be delegated through permissions on the GPO.
- GPOs cannot be linked to the Users and Computers containers in AD. This is one important reason it is recommended that users and computers be moved into OUs.

Considering the impact of GPO inheritance and the support for delegation of GPO administration, it becomes apparent that intelligent OU design can be quite important. If OU design is carefully considered before domain implementation, it can save considerable time later.

#### **C.2.5.4. Group Policy Management Console (GPMC) and Group Policy Results Tool**

There are multiple ways for an administrator to maintain and evaluate GPO application, but two tools deserve particular mention: the GPMC and the Group Policy Results tool.

The GPMC consists of an MMC snap-in and some scriptable interfaces for managing Group Policy. The GPMC can be used from Windows Server 2003 or Windows XP computers and it can be used to manage Group Policy on Windows 2000 Server computers as well as Windows Server 2003 computers. The GPMC package can perform editing, backup/restore, and import tasks. GPOs can be copied from one domain to another by using the GPMC backup function in the source domain and the import function in the target domain.

The GPMC provides the Group Policy Modeling interface, formerly known as Resultant Set of Policy (RSoP) planning mode, and the Group Policy Results interface, known formerly as RSoP logging mode. The GPMC is the recommended tool for integrated management of Group Policy.

The Group Policy Results Tool (gpresult.exe) is a command line tool that can be used to display the effective group policy for the current user and computer that results from the application of all GPOs at the local, site, domain, and OU levels. This can be an effective tool for diagnosing GPO application issues.

#### **C.2.5.5. GPO Auditing and Backup**

The final Group Policy subjects are auditing and backup. Auditing changes to GPOs can be done in much the same way as auditing changes to data files. Each GPO has audit settings that can be enabled. For high security environments such as the DoD, proper values for audit settings are required by policy, important for incident investigation, and necessary to ensure the capture of data required for prosecution of malicious behavior.

There are multiple ways to back up GPOs. The GPMC can be used to save GPOs for recovery and for distribution to other domains. GPO data is also backed up when the Windows System State data is backed up. This was briefly discussed in Section C.2.2, Forest and Domain Architecture. The use of the GPMC is highly recommended over the System State backup method because of the granularity of the GPMC operation. In environments with active use of GPOs, the GPMC tool is a more efficient approach to the task and could significantly reduce recovery time.

#### **C.2.6. Ports and Protocols**

It is noted in Section 2.2.5, Network Ports and Protocols, that directory servers provide and consume services that involve network traffic. Because of the close integration of AD with Windows, it is difficult to identify those ports and services that are unique to the function of AD. The objective of this section is to identify the network services, standard ports, and the associated AD and Windows services that are needed for an AD environment.

AD is a network service provider in two respects:

- On each domain controller, AD responds to directory query and update traffic using LDAP protocol on port 389. When LDAP data-signing is required on the domain controller (as enforced in the DoD Windows Checklists), the integrity of the transmission is protected. Clients can invoke SASL-based encryption for confidentiality. When a properly configured PKI certificate is installed, AD supports encrypted sessions to respond to directory query and update traffic using LDAPS protocol on port 636.
- On domain controllers that are also GC servers, AD responds to directory query traffic in LDAP protocol on port 3268. As with the port 389 traffic, integrity and confidentiality protection is supported. When a properly configured PKI certificate is installed, AD supports encrypted sessions to respond to directory query traffic using LDAPS protocol on port 3269.

AD and the Windows services that AD depends on are network service consumers in several respects:

- Operations such as configuring access permissions can use an LDAP query to a domain controller on port 389 to enumerate user and group accounts.
- When a computer or user logs on to a domain controller, that domain controller contacts a GC server on port 3268 or 3269 to determine universal group membership.
- AD accesses DNS servers on port 53 for multiple purposes. Domain controllers update DNS service location records that are read by AD clients. Domain controllers query DNS servers to obtain the addresses of other Windows servers.
- AD uses ICMP ping packets between domain controllers and AD clients to evaluate link speed for GPO processing.
- The Windows FRS uses the Microsoft-DS service on port 445 for AD data replication and the RPC Endpoint Mapper service on port 135.
- The domain controller that holds the PDC Emulator FSMO role can provide a time synchronization service using NTP on port 123.
- The Windows Local Security Authority (LSA) works in conjunction with AD for authentication and authorization. It uses several ports. When Kerberos authentication and authorization are utilized, port 88 is used.

The following table lists the network services, ports, and related Windows services that are required for AD functions. Although all of these ports are not used with every domain controller and client, they would be used somewhere in a complete AD configuration.

Port	TCP\UDP	Service	Windows System Service
N/A	N/A	ICMP (ping)	Group Policy
53	TCP\UDP	DNS	DNS
88	TCP\UDP	Kerberos	Key Distribution Center
123	UDP	NTP \ SNTP	Windows Time
135	TCP	RPC Endpoint Mapper	LSA, FRS, Group Policy
389	TCP\UDP	LDAP	LSA, Group Policy
445	TCP	Microsoft-DS / SMB	Group Policy
636	TCP\UDP	LDAPS	LSA
3268	TCP	MS Global Catalog	LSA
3269	TCP	MS Global Catalog SSL	LSA
1024-65536	TCP	RPC (dynamic)	LSA, FRS, Group Policy

**Table C-6. AD Port\Protocol Use**

Please note that this information is subject to change as Microsoft's implementation changes. Although the data is presently consistent with the information in Microsoft Knowledge Base article 832017, users should check the Microsoft documentation before changing configurations based on this information.

It must also be noted that although almost all of these Windows services work only on the documented port numbers, some of these services can be configured to use different port numbers. By default the LSA and FRS Windows services support RPC network service calls on a dynamic port number above 1023. But it is possible to create Windows registry entries that override that behavior and specify explicit, fixed values.

- In key HKLM\SYSTEM\CurrentControlSet\Services\NTDS\Parameters, the “TCP/IP Port” value can be used to specify a port for AD replication traffic.
- In key HKLM\SYSTEM\CurrentControlSet\Services\NTFRS\Parameters, the “RPC TCP/IP Port Assignment” value can be used to specify a port for FRS traffic.

Refer to the Microsoft documentation for details. Please note that changes to these parameters must be done with extreme care and only in coordination with local network management personnel. Errors in registry settings or a firewall block of a specific port could result in the failure of services that support AD.

Directory synchronization tools that can be used with AD also rely on access to network services. These tools are very briefly described in Section 2.3, Directory Synchronization Tools and Technology. The following table lists the network services, ports, and related tool or technology.

Port	TCP\UDP	Service	Tool\Technology
53	TCP\UDP	DNS	MIIS
80	TCP	HTTP	MIIS using DSfW
88	TCP\UDP	Kerberos	MIIS
135	TCP	RPC Endpoint Mapper	ADAM
389	TCP\UDP	LDAP	ADAM, SimpleSync, MIIS
443	TCP	HTTPS	MIIS using DSfW
464	UDP	Kerberos kpasswd	MIIS
636	TCP\UDP	LDAPS	ADAM, SimpleSync, MIIS

**Table C-7. Synchronization Port\Protocol Use**

As with the information in the AD Port\Protocol Use table, users should check the Microsoft documentation before using this data.

Some of the network services in these lists have been associated with significant security vulnerabilities. In particular, some services associated with ports 135 and 445 are identified as high risk. For this reason, their use across certain network boundaries is restricted. DoD policy on this is described in *DoD Instruction 8551.1, Ports, Protocols, and Services Management (PPSM)* and the associated *Ports, Protocols, and Services (PPS) Assurance Category Assignments List*. The *PPS Assurance Category Assignments List* provides detailed guidance for specific services and port numbers. It is available to DoD and Government users through the IASE web site.

In order to support an AD environment that spans DoD enclave boundaries and is compliant with the DoD policy, the use of encrypting VPN technology is the only known acceptable solution. Please refer to the *Network Infrastructure STIG* for information on VPN implementation.

### **C.3. Technology-Specific Security Requirements**

This section supplements the information in Section 3, Directory Service Security Requirements. It provides guidance on meeting the general requirements, additional requirements, and recommendations that are unique to AD as a directory service.

It is important to understand specific terminology used in this section:

- AD database - AD data is stored primarily in a file named ntds.dit. This is not a database in the sense of a general-purpose data store, but it does provide the data repository for AD and so is referred to by the term AD database in this document.

It should be noted that the requirements here are based on the versions of Microsoft Windows 2000 Server and Windows Server 2003 with the current service packs and security fixes at the time this appendix was written. Specifically, changes introduced with Windows Server 2003 Release 2 are not reflected. As with the implementation of all security configuration guidance, DoD Components should test configuration settings in a test environment before implementation in production to ensure their specific environment is not impacted in unintended ways.

#### **C.3.1 Security Design and Configuration**

Section 3.1, Security Design and Configuration, discusses five requirement areas related to the IA Controls in the Security Design and Configuration subject area. A few of these requirement areas have unique considerations for AD. The following sections provide details and guidance on meeting the general requirements for AD security in these areas.

Section 3.1.2, Configuration and Implementation Integrity, referred to cross-directory authentication mechanisms. Within AD, certain types of AD trust relationships constitute cross-directory authentication mechanisms.

Most manual AD trusts are defined interconnections between two or more parties in different forests. While a trust facilitates resource sharing, it weakens the isolation provided by the AD forest architecture. It is noted in Section C.2.4, Trust Relationships, that external, forest, and realm AD trusts enable a user from a trusted domain or forest to access resources in the trusting domain or forest without requiring the re-authentication of that user. If a trust were defined improperly or between the wrong parties, it could allow access to resources that had weak access control permissions.

To comply with the requirement (DS00.1120) in section 3.1.2 to collect documentation for cross-directory authentication configurations, the following information has to be maintained for each external, forest, and realm AD trust defined on each forest or domain:

- Trust type (external, forest, or realm)

- Fully qualified domain names of each party
- MAC, confidentiality, and classification levels of each party
- Direction (one-way or two-way)
- Transitivity
- Status of selective authentication option and SID filtering (quarantine) option.

Section 3.1.4, Software Integrity, described the need to secure access to programs that manipulate directory data. There are many instances of such programs in the AD environment. Windows server software comes with several command-line and GUI programs that can display, add, modify, and delete AD data as well as configure the AD environment. Some of these programs are:

- Multi-purpose - The AD Installation Wizard (dcpromo.exe), NTDS Utility (ntsdutil.exe), and Group Policy Restore Utility (dsgpofix.exe) programs perform various functions.
- Import and export - The CSVDE (csvde.exe) and LDIFDE (ldifde.exe) programs can import to and export from the AD database.
- MMC snap-ins - Various administrative functions are scattered across the AD Users and Computers (dsa.msc), AD Domains and Trusts (domain.msc), AD Sites and Services (dssite.msc), AD Schema (schmmgmt.msc), Group Policy (gpedit.msc), Domain Controller Security Policy (dcpol.msc), and Domain Security Policy (dcompol.msc) interfaces.
- Windows Server 2003 additions - The directory service add (dsadd.exe), display (dsget.exe), modify (dsmod.exe), move (dsmove.exe), query (dsquery.exe), and remove (dsrm.exe) programs were added in Windows Server 2003 to provide additional command-line AD data manipulation tools.

These programs are considered sensitive because they support administrative updates to AD or they support the extraction of various scopes of information from the AD database. If an unauthorized user is allowed to access these programs that is one step in allowing that user to compromise the confidentiality, integrity, or availability of the AD environment or data.

Existing requirements in the *Windows 2003/XP/2000/Vista Addendum* address access control for the programs installed by default with the server software.

Microsoft also makes programs available on the Windows server CD and through download from their web site that can manipulate AD data and the environment. Some of these programs are:

- The Domain Controller Diagnostic tool (dcdiag.exe) analyzes the Windows environment for issues that affect AD.
- The Directory Service Access Control List tool (dsacls.exe) allows display and modification of AD object permissions.
- The LDAP utility (ldp.exe) is a general tool to display and modify AD data.
- The Domain Manager tool (netdom.exe) performs a variety of functions including creation of trusts and management of domain computer accounts.

- The Domain Secure Channel utility (nltest.exe) can be used to display and test AD trust relationships.
- The Replication Administration (repadmin.exe) and Replication Monitor (replmon.exe) tools display and modify AD replication functions.
- The ADSI Edit MMC snap-in (adsiedit.msc) allows detailed display and modification of AD database objects.

These additional tools are installed as Support Tools in Windows. As with the Microsoft programs identified earlier, these tools allow administrative updates to AD or they allow various scopes of information to be extracted from the AD database. If an unauthorized user is allowed to access these programs, that is one step in allowing that user to compromise the confidentiality, integrity, or availability of the AD environment or data.

Depending on the environment, the default installation of these programs does not provide adequate access control. To comply with the requirement (DS00.1150) in section 3.1.4, it is necessary to configure the permission specifications included here.

Object	Name	Type	Access
...\%ProgramFiles%\Support Tools\	Administrators SYSTEM [Other IAO-authorized groups]	Allow Allow Allow	Full Control Full Control Read, Execute  With propagation

**Table C-8. Support Tools Access Permissions**

### C.3.2 Identification and Authentication

Section 3.2, Identification and Authentication, describes general requirements for identification and authentication data protection for directory services. This section identifies the account and authentication data requirements unique to AD.

The DSRM system state used for AD restore operations was mentioned in Section C.2.2, Forest and Domain Architecture. When a Windows server is promoted to be a domain controller, a password must be selected for use when that domain controller is booted into DSRM. In DSRM, the user at the server console has the ability to manipulate the AD database file without the software-enforced protections of the normal domain security environment. If a malicious user gained physical access to a domain controller and obtained this password, he would be able to capture, modify, or delete the AD database. Depending on the AD object, malicious modifications might later be propagated throughout the forest by normal AD replication.

To strengthen control of access to DSRM, password complexity requirements must be followed.

- *(DS10.0150: CAT II) The IAO will ensure the password defined for use in DSRM complies with the password content requirements (length and composition) as specified in the current DoD instructions and JTF-GNO documents.*

The password that controls DSRM access resides in the SAM file on the domain controller. When the domain controller is running in the normal Active Directory mode, the SAM file is not accessible to the standard tools that check for the presence and age of passwords.

Password expiration is required by current DoD security policy. In the case of the DSRM-related password, periodic changes enhance security by ensuring that a password is defined, the correct value is authoritatively known, and a potentially compromised password is replaced. Because access to the password is very limited, the yearly expiration period specified in the *Windows 2003/XP/2000/Vista Addendum* for application accounts is appropriate.

- *(DS10.0151: CAT II) The IAM will ensure there is a local policy in place that requires the password defined for use in DSRM to be changed on a yearly basis.*

There are two methods available to change this password. The first requires a service outage for the domain controller. The server is booted into DSRM, a “net user” command or Local Users and Groups MMC snap-in is invoked to change the password, and the server is booted again to return to AD mode. The second method can be done without a service interruption. In Windows 2000 Server, the setpwd utility can be used. In Windows Server 2003, the “set dsrm password” subcommand of the NTDS Utility (ntsdutil.exe) can be used. Please refer to the Microsoft documentation for specific command information.

A requirement to restrict physical access to the DSRM password is stated in Section C.3.5, Continuity.

### **C.3.3 Enclave and Computing Environment**

Section 3.3, Enclave and Computing Environment, discusses eight requirement areas related to the IA Controls in the Enclave and Computing Environment subject area. Several of these requirement areas have unique considerations for AD. The following sections provide guidance on meeting the general requirements, additional requirements, and recommendations that are unique to AD security in these areas.

#### **C.3.3.1. Architecture and Cross-Directory Authentication (AD Trusts)**

Section 3.3.2, Architecture and Cross-Directory Authentication, provides a very general discussion of the impact of directory implementation architecture and cross-directory authentication. This section identifies the unique aspects of those issues for AD.

The way in which AD is implemented in terms of domain and forest architecture has an enormous impact to Windows security support of user isolation. The type and characteristics of manually defined AD trusts, as an implementation of cross-directory authentication, also has a very significant impact. The resulting AD environment profoundly affects account and resource definitions, user authentication, and resource access control.

The subject of AD architecture and trusts is complex. Please reference the information in Section C.2.2, Forest and Domain Architecture, and section C.2.4, Trust Relationships, for a brief technology background on these subjects.



The large number of individual needs and environmental constraints within an organization the size of DoD makes it impossible in a practical sense to define a comprehensive list of strict AD architecture security requirements. A configuration that provides an appropriate level of security in one instance could be unreasonably weak in another. As a result, most of the guidance here is written as recommendations that the Components need to interpret for their particular environment.

The following guidance should be applied when determining when to implement unique forests or domains:

- The IAO should ensure separate AD forests are implemented when there is a need for one group of users and resources to be isolated from another. Isolation refers to two considerations: trusts and administration. Whereas domains in a single forest are automatically connected by transitive, two-way trusts, there are no default trust relationships between forests. The second issue is that administrators of one forest are not able to change the security configuration of the users and resources in another forest.

One example in which separate forests are recommended is a perimeter configuration such as a DMZ. If external users accessing DMZ resources have no requirement to access resources on the internal side of the DMZ, a separate forest for the DMZ should be considered.

- The IAO should ensure separate AD domains are implemented when there is a need for one group of users to work autonomously from another. Autonomy also has trust and administration considerations. Because domains in a forest are automatically connected by transitive, two-way trusts, users can be granted access to resources in other domains without the need for additional user accounts. From the administration perspective, domain administrators can perform all configuration tasks for the domain, but some administrators of the root domain still have the ability to change the security configuration of the users and resources in a child domain.
- The IAO should consider the use of an empty root domain when the number of domains in the forest is large and specific requirements dictate the need for individual domain administration.

In this case stronger security policies (such as requirements for longer passwords) can be implemented to protect the more sensitive forest-wide roles used in the root domain. Fewer users can be defined in the root domain, thus reducing the attack surface of that domain. And forest root domain controllers can be dedicated to their AD functions.

- The IAO should ensure a separate AD domain is implemented for a software development project when the activities of that project have a high potential to impact a production environment.

It must be acknowledged that AD implementations consisting of the fewest forests and domains generally provide ease of administration, less complexity, and lower cost. However, it must also be understood that fewer security barriers accompany those benefits. Such barriers might deter or prevent a successful attack by an insider or an intruder who compromises a domain account.

One measure that can be taken to improve the implemented AD architecture is the exclusion of Windows NT servers as domain controllers. Windows NT does not support the same level of security that can be achieved with Windows 2000 Server or Windows Server 2003. Also, because Windows NT is no longer supported, there are un-patched vulnerabilities that increase the risk of successful attack. Finally, when a Windows NT server is a domain controller, it forces Windows 2000 Server or Windows Server 2003 servers in the same domain to be configured less securely in order to support interoperability.

It must be noted that the *Windows 2003/XP/2000/Vista Addendum* designates the use of unsupported software such as Windows NT Server, as a domain controller or for any other purpose, to be a Severity Category I finding. This reflects the fact that Microsoft no longer provides support or patches for Windows NT Server.

In the past operational requirements may have precluded a migration from Windows NT. However, in many cases migrations are now complete and Windows NT is no longer used as a domain controller. In these environments, preventing the addition of new Windows NT domain controllers does enhance security. This can be accomplished by setting the domain functional level to a value that does not allow a Windows NT domain controller.

- *(DS10.0160: CAT III) In AD domains that have no Windows NT domain controllers, the IAO will ensure the domain functional level is Windows 2000 native or Windows Server 2003.*

As with AD architecture, it is impossible in a practical sense to define a comprehensive list of strict AD trust security requirements for all DoD environments. Following are a few requirements and some recommendations to develop a more secure environment.

As noted in Section C.2.4, Trust Relationships, AD trusts are defined automatically between domains in a forest and can be defined manually between forests and between domains in different forests. If a trust is defined between two domains and the permissions for a file in the trusting domain do not sufficiently restrict access, the contents of the file could be disclosed or modified by an unauthorized user in the trusted domain. Because trusts eliminate a level of authentication, it is very important that they are defined only when needed.

- *(DS10.0170: CAT II) The IAO will ensure external, forest, and realm AD trusts are defined only as required to support authorized access.*

Configuring an AD trust between systems that are at different DoD classification levels could substantially increase risk. If desired, such a configuration would require the use of a controlled interface and it has not been determined that such a configuration for AD trusts could be adequately secured. Depending on the direction of the trust, data at one classification might be made available to a user in a lower classification level.

- A trust could allow a user in the trusted domain at a lower classification level to read data in the trusting domain at a higher classification level.
  - A trust could allow a user in a trusted domain at a higher classification level to copy data to the trusting domain at a lower classification level.
- *(DS10.0180: CAT I) The IAO will ensure external, forest, and realm AD trusts are not configured between systems at different classification levels.*

Configuring an AD trust between a DoD system and a non-DoD system could also substantially increase risk. A compromise in the non-DoD forest or domain could lead to a compromise of the DoD forest or domain. However, there may be cases in which information sharing requirements between DoD and non-DoD Government agencies, coalition partners, or contractors justify such a configuration. In order to ensure the risks of this configuration are properly recognized and mitigating actions for the network connections are taken, there are multiple conditions for this configuration.

- *(DS10.0181: CAT I) The IAO will ensure external, forest, and realm AD trusts are not configured between DoD and non-DoD systems unless:*
  - *The network connections comply with all requirements for external connections defined in the Network Infrastructure STIG, including a Memorandum of Agreement (MOA) between the two parties*
  - *Explicit approval of the trust by the DAA is documented.*

It is noted in Section C.2.4.3, Manually Defined Trusts, that the function that enables the SID History feature in Windows could be a vulnerability leading to an elevation of privilege attack. If a user in a trusted domain were able to forge credentials passed to a trusting domain, that user might gain unauthorized, privileged access to the resources in the trusting domain. In order to help prevent the use of forged credentials over trusts, the SID filtering option must be enabled.

- *(DS10.0190: CAT II) The IAO will ensure SID filtering is enabled on all external and forest trusts.*

**NOTE:** Implementation of this requirement could have a significant impact. Although SID filtering is the default for trusts created under Windows 2000 Server with SP4 (and later) and Windows Server 2003, the setting can be altered. Without proper review and update of resource permissions to grant authorized access, implementation of the requirement could result in denied access to authorized users.

It should be noted that enabling SID filtering could have an impact on the use of Universal groups. A user in a trusted domain who was a member of Universal group in a different domain, would not be able to access resources in the trusting domain that were secured by access permissions using the Universal group. In this case a new group should be created in the trusted domain and the access permissions of the resource in the trusting domain should be updated.

Although the use of SID History is convenient, it should generally be used as a short-term measure during user migrations between domains. Removing SID History values that are no longer used enhances inter-domain security.

- The SA should periodically (at least semiannually) review the AD database and clear sIDHistory attribute values that are no longer needed.

With the addition of forest trusts in Windows Server 2003, it is possible to simplify trust configurations. A single forest trust could replace a large number of individual external trusts. However, because a forest trust may be effective over a much wider scope, an additional access control is needed to maintain a high security environment. The Selective Authentication trust option and the Allowed to Authenticate permission provide and enable this control. If these items are not properly configured, a user in a trusted forest might be able to gain unauthorized access to a resource with weak access permissions in the trusting forest.

When the Selective Authentication option is set on a forest trust, the Allowed to Authenticate permission must be configured on the resource server. The computer in the trusting forest in which the resource is located must have the Allowed to Authenticate permission granted to the user or group in the trusted forest that wants to access the resource. In this way access through a forest trust is granted only to the users specifically authorized by the administrator of the server in the trusting forest.

- *(DS10.0200: CAT II) The IAO will ensure the Selective Authentication option is enabled on all forest trusts.*

**NOTE:** Implementation of this requirement could have a significant impact. Without proper review and update of the Allowed to Authenticate permission, implementation of the requirement could result in denied access to authorized users.

The following additional guidance should be applied when determining how to implement AD trusts.

- The IAO should ensure whenever possible one-way trusts are defined instead of two-way trusts or trusts in both directions.

A domain in a perimeter configuration defined in a separate forest could be a good candidate for a one-way trust. In this case the trust would be defined with an internal domain as the trusted domain and the perimeter domain as the trusting domain. This would allow users authenticated on the internal domain to update resources in the perimeter domain without separate credentials or re-authentication.

- When there is a need for a large majority of domains in one forest to trust the domains in another forest and the risk has been determined to be acceptable, the IAO should ensure a forest trust is defined instead of multiple external trusts.
- When a forest spans a wide area network, the IAO should ensure shortcut trusts are defined. This can improve intra-forest trust processing and help to avoid availability problems that could be caused by network outages along the trust path.

### C.3.3.2. Data Access Control - AD Files

Section 3.3.3, Data Access Control - Files, describes the general requirement for access control for directory service data files. This section identifies the unique file system objects for AD.

AD-specific data files include the AD database, log files, work files, and the GPT directory (SYSVOL). FRS, a Windows component critical to supporting AD replication, stores data in a database and log files. If these elements are located in the default locations and the guidance in the *Windows 2003/XP/2000/Vista Addendum* is followed, inherited permissions would provide adequate access control. But because this data can be moved elsewhere and it is critical in order for AD to function, additional attention to the permissions is needed. To comply with the requirement (DS00.0120) in section 3.3.3, it is necessary to configure the permissions included here.

Please note that the permissions for the database, log, and work files differ between Windows 2000 Server and Windows Server 2003. The permissions for the account names with an asterisk in the following table are only required for Windows Server 2003.

Component	Object	Name	Type	Access
Database	...\ntds.dit	Administrators SYSTEM CREATOR OWNER* Local Service*	Allow Allow Allow	Full Control Full Control [None on file] Create Folders / Append Data
Log files and log reserve files	...\edb*.log, ...\res1.log ...\res2.log	Administrators SYSTEM CREATOR OWNER* Local Service*	Allow Allow Allow	Full Control Full Control [None on file] Create Folders / Append Data
Work files	...\temp.edb ...\edb.chk	Administrators SYSTEM CREATOR OWNER* Local Service*	Allow Allow Allow	Full Control Full Control [None on file] Create Folders / Append Data

Component	Object	Name	Type	Access
GPT parent directory	...\SYSVOL	Administrators	Allow	Full Control Read, Read & Execute, List Folder Contents [None on dir.] Read, Read & Execute, List Folder Contents Full Control
		Authenticated Users	Allow	
		CREATOR OWNER Server Operators	Allow	
GPT policies directory	...\SYSVOL\ domain\Policies	Administrators	Allow	Full Control Read, Read & Execute, List Folder Contents [None on dir.] Read, Read & Execute, List Folder Contents, Modify, Write Read, Read & Execute, List Folder Contents Full Control
		Authenticated Users	Allow	
		CREATOR OWNER Group Policy Creator Owners	Allow	
FRS directory	...\Ntfrs	Server Operators	Allow	Full Control Full Control
		SYSTEM	Allow	
		Administrators	Allow	
		SYSTEM	Allow	

**Table C-9. AD Data Access Permissions**

It should be noted that the full path to these files has to be determined from the Windows registry entries for AD. This is due to the fact that other products utilizing Microsoft's Extensible Storage Engine create files with some of the same names.

### C.3.3.3. Data Access Control - AD Database Objects

Section 3.3.4, Data Access Control - Directory Database Objects, describes the general requirement for access control for directory database objects. This section identifies specific database objects with required access settings and additional unique requirements for AD.

Although the access permissions for all AD objects are important, there are two types that require more vigilant review because they are likely to be altered during normal administrative activities. These types are GPOs and OUs. The following table lists the object access permissions that are required for compliance with the general requirement (DS00.0130) in section 3.3.4.

Object	Name	Type	Access
[Group Policy - e.g., Default Domain]	Administrators	Allow	Full Control
	Creator Owner	Allow	Full Control
	SYSTEM	Allow	Full Control
	ENTERPRISE DOMAIN CONTROLLERS*	Allow	Read
	Authenticated Users [or other user groups]	Allow	Read
			Apply Group Policy
[Organizational Unit - e.g., Domain Controllers]	Administrators	Allow	Full Control
	Creator Owner	Allow	Full Control
	SYSTEM	Allow	Full Control
	Authenticated Users [or other user groups]	Allow	Read

**Table C-10. AD Database Object Access Permissions**

Please note that different permissions may be required to implement specific Group Policy designs and delegated administration. These variations are acceptable when documented by the IAO:

- It is anticipated that the Apply Group Policy permission could be set to Deny in some cases as part of the exclusion of a specific group policy using security filtering.
- When OU administration is delegated, permissions beyond Read may be allowed to groups authorized to administer those OUs.

As ongoing research is completed, it is expected that guidance for access permissions for other AD objects will be provided in future versions of this document.

As a mechanism to enable bulk access to AD objects, the “Synchronize directory service data” user right was defined for Windows domain controllers. An account granted this right is allowed to read all AD objects and properties, bypassing the access control permissions defined for them.

The high-level privilege associated with this right is too powerful for any normal environment. If an unauthorized user has access to an account with this right, the confidentiality of all data stored in the AD database is compromised. Since Windows processes that need this level of access execute under the SYSTEM account, all AD objects are defined with permissions that allow access by SYSTEM, and granular access permissions can be defined for individual objects, there is no known need to grant the Synchronize directory service data user right.

- (DS10.0210: CAT I) The IAO will ensure no accounts are granted the Synchronize directory service data user right.

AD was initially implemented to allow anonymous access to object data. This was primarily done to provide backward compatibility to older AD clients. Some early Windows programs, most notably functions implemented in Windows NT, were written to access AD object data

through anonymous connections. Consequently when this access is disabled in AD, certain Windows NT functions no longer work.

Since access to AD objects is controlled through ACLs, anonymous access can be enabled by defining object ACLs that permit access to groups that contain anonymous users. Microsoft implemented this through a built-in Windows group named "Pre-Windows 2000 Compatible Access". If the Everyone group is nested in the Pre-Windows 2000 Compatible Access group and anonymous users are part of the Everyone group, anonymous access to AD data is permitted.

In Windows Server 2003 an additional control over anonymous access to AD was implemented. This control is through the dsHeuristics attribute of the Directory Service object. The default setting of this attribute specifies that only authenticated users may initiate an LDAP request. Because the Directory Service object is in the Configuration partition of the AD database, and that partition is replicated forest-wide, the setting of the dsHeuristics attribute is effective for all Windows Server 2003 domain controllers in a forest.

In environments using current Windows software, there should be no need to allow anonymous access to AD data. If an unauthorized user gains anonymous access, that data could be useful in subsequent attacks on the forest or domain. While some of the data would be insignificant, there are elements such as user and group names that are very significant for devising an attack and selecting a target. For this reason, AD settings must be configured to prevent anonymous access.

- *(DS10.0220: CAT II) The IAO will ensure the Pre-Windows 2000 Compatible Access group does not contain the Everyone or Anonymous Logon groups.*
- *(DS10.0230: CAT II) The SA will ensure the dsHeuristics attribute is configured to prevent anonymous access.*

**NOTE:** Configuration of the dsHeuristics attribute to disable anonymous access is effective only on domain controllers running Windows Server 2003.

**NOTE:** Implementation of these requirements could have a significant impact. In domains including Windows NT hosts, implementation of the requirement could result in denied access to authorized users and processes running on Windows NT hosts.

Beyond the specific requirement for the Everyone group, the following recommendation applies:

- The IAO should restrict the membership of the Pre-Windows 2000 Compatible Access group to few or no members.

#### **C.3.3.4. Data Change Auditing - AD**

Section 3.3.5, Data Change Auditing, describes the general requirement for auditing changes to directory data. This section identifies specific database objects with required unique audit settings for AD.



To track changes to the AD database, log files, work files, and the GPT directory (SYSVOL), specific audit settings for the files are necessary. If the guidance in the *Windows 2003/XP/2000/Vista Addendum* and the *Windows Server 2003 Security Guide* is followed, there is no need for additional audit setting requirements for data files.

Because of their importance to the operation of AD, special attention is necessary for objects in the domain partition of the AD database. To ensure changes and attempted changes to the sensitive domain partition objects are tracked, the following are the minimum audit settings that are required for compliance with the general requirement (DS00.0140) in section 3.3.5.

Please note the following for these items:

- “*Domain*” is the actual Windows domain name in LDAP form. For example, this could be “DC=aofn21,DC=disa,DC=mil”.
- The object names are expressed as they would be seen in an AD directory tree display. This contrasts with a typical LDAP string in which the Relative Distinguished Names (RDNs) would be sequenced from the most specific to least specific, as in “CN=AdminSDHolder, CN=System, *Domain*”.

<b><i>Domain Object</i></b>			
<b>Type</b>	<b>Name</b>	<b>Access</b>	<b>Scope</b>
Failure	Everyone	[All access types]	<i>Domain object only</i>
Success	Everyone	Write All Properties Modify Permissions Modify Owner	<i>Domain object only</i>
Success	Administrators	All Extended Rights	<i>Domain object only</i>
Success	Domain Users	All Extended Rights	<i>Domain object only</i>

<b><i>Domain,CN=System,CN=AdminSDHolder Object</i></b>			
<b>Type</b>	<b>Name</b>	<b>Access</b>	<b>Scope</b>
Failure	Everyone	[All access types]	AdminSDHolder object only
Success	Everyone	Modify Permissions Modify Owner Write All Properties	AdminSDHolder object only

<b><i>Domain,CN=System,CN=Policies Object and Domain,CN=System,CN=Policies,CN=GPC Child Objects</i></b>			
<b>Type</b>	<b>Name</b>	<b>Access</b>	<b>Scope</b>
Failure	Everyone	[All access types]	Policies object and all <i>GPC child</i> objects

<b><i>Domain,CN=System,CN=Policies Object and Domain,CN=System,CN=Policies,CN=GPC Child Objects</i></b>			
<b>Type</b>	<b>Name</b>	<b>Access</b>	<b>Scope</b>
Success	Everyone	Modify Permissions Modify Owner Create groupPolicyContainer Objects Delete Delete groupPolicyContainer Objects Delete Subtree	Policies object only
Success	Everyone	Modify Permissions Write All Properties	all <i>GPC child</i> objects

<b><i>Domain,CN=System,CN=RID Manager\$ Object</i></b>			
<b>Type</b>	<b>Name</b>	<b>Access</b>	<b>Scope</b>
Failure	Everyone	[All access types]	RID Manager\$ object only
Success	Everyone	All Extended Rights Write All Properties	RID Manager\$ object only

<b><i>Domain,CN=Infrastructure Object</i></b>			
<b>Type</b>	<b>Name</b>	<b>Access</b>	<b>Scope</b>
Failure	Everyone	[All access types]	Infrastructure object only
Success	Everyone	All Extended Rights Write All Properties	Infrastructure object only

<b><i>Domain,OU=Domain Controllers Object</i></b>			
<b>Type</b>	<b>Name</b>	<b>Access</b>	<b>Scope</b>
Failure	Everyone	[All access types]	Domain Controllers OU and all child objects
Success	Everyone	Modify Permissions Modify Owner Create All Child Objects Delete Delete All Child Objects Delete Subtree	Domain Controllers OU only
Success	Everyone	Write All Properties	Domain Controllers OU and all child objects

**Table C-11. Domain Partition Object Audit Settings**

As ongoing research is completed, it is expected that guidance for audit settings of other AD objects will be provided in future versions of this document.

Because the impact of these settings depends on activity in the local environment, SAs are advised to monitor the Security event logs on their domain controllers and make adjustments for any increase in the amount of log data.

### **C.3.3.5. Group Membership and Limiting Privileges - AD**

Section 3.3.6, Group Membership and Limiting Privileges, describes general requirements for limiting membership in groups with privileged access to the directory. This section identifies additional specific group and privilege control requirements and recommendations that are unique for AD.

It is noted in Section C.2.3, Group Membership, that several groups are automatically created when Windows is installed that have default privileges over AD files and database objects. Membership in some of these groups grants permissions to the members to change configuration settings that can impact an entire AD forest. With the exception of the Domain Admins group, membership in the other groups is generally not required to support common administrative tasks.

If an unauthorized user is able to gain membership in the Enterprise Admins, Schema Admins, or Incoming Forest Trust Builders groups, that user would be able to display, add, or change AD objects that could impact the confidentiality, integrity, or availability of an entire forest. Gaining membership in the Domain Admins or Group Policy Creator Owners groups would allow the user to display, add, or change AD objects that could impact the confidentiality, integrity, or availability of an entire domain.

- *(DS10.0240: CAT II) The IAM will limit the number of users and document those users assigned to the following security groups: Domain Admins, Enterprise Admins, Schema Admins, Group Policy Creator Owners, and Incoming Forest Trust Builders.*

Because of the individual needs and environmental constraints within DoD Components, it is impossible in a practical sense to define a comprehensive list of strict group membership requirements. A hierarchy that provides appropriate control for a large unit could be unreasonably complex in another. As a result, most of the guidance here is written as recommendations that the Components need to interpret for their particular environment.

The following guidance should be applied when managing group membership:

- The IAO should ensure the number of users assigned to the Domain Admins group of an empty or small forest root domain is limited to a very small number of users.
- The IAO should consider restricting membership in the Enterprise Admins, Schema Admins, and Incoming Forest Trust Builders groups to no (zero) members and adding a local Administrator account to the groups temporarily when needed.

Because membership in Schema Admins would only be required when updates to the AD schema are needed, such as the installation of products like Microsoft Exchange or Systems Management Server, the Schema Admins group should almost always be empty.

- The IAO should consider the use of group policy that uses the Restricted Groups setting. Implementing such a policy would provide periodic, automated refreshes of the memberships of privileged groups.

Although the concern is not confined to AD or privileged users, the structure of Windows groups deserves a brief discussion because it affects, and is affected by, AD implementation. The various group types that can be used in Windows were briefly discussed in Section C.2.3, Group Membership. The assignment of privileges and access permissions to groups rather than individual users is done in almost all environments as an implementation of RBAC that simplifies security administration.

It is noted previously that the domain functional level of a domain affects the group implementation. Most notably it occurs when domains are raised to the Windows 2000 native domain functional level. There are two items of significant impact to group structure.

- Universal groups are available. Their membership is replicated to all Global Catalog servers.
- The ability to nest domain global groups within domain global groups is available.

The structure of groups must be designed for the specific environment. The following guidance is recommended:

- The IAO should ensure access and privilege permissions are assigned to domain local groups rather than individual accounts or other types of groups. This helps to preserve control of permissions at the domain level.
- The IAO should ensure Universal groups are used to define groups that span domains and that single-domain forests do not use Universal groups. When Universal groups are not used, dependence on the Global Catalog server is reduced.
- The IAO should ensure domain global groups rather than individual accounts are members of Universal groups. In addition to being more efficient because it reduces replication traffic among Group Catalog servers, it helps to preserve control of user memberships at the domain level.

As a reminder, these recommendations can be thought of as implementations of the A-G-DL-P, A-G-G-DL-P, and the A-G-G-U-DL-P strategies that were discussed in Section C.2.3, Group Membership.

As noted in Section 3.3.6, Group Membership and Limiting Privileges, it is possible to limit the scope of privileges assigned to accounts and groups by organizing directory objects and using delegation mechanisms. The creation and delegation of OUs for users and computers within AD is an example.

Creating an OU for user accounts that have similar privileges and other OUs for users without can provide groupings to which appropriate group policy can be linked. A similar strategy can be applied to computers. Just as domain controller computer accounts are assigned to the Domain Controllers OU and that OU has the Default Domain Controllers Policy linked to it, it is

desirable to create OUs for other categories of computers and create and link policies to them. Once OUs have been created it is also possible to delegate control over those OUs. This allows the creation of privileged users who are limited in scope.

The structure of OUs and the delegation strategy must be designed for their specific environment. The following guidance is recommended:

- The IAO should ensure user accounts and computer accounts are assigned to OUs.
- The IAO should consider the use of group policies to limit access to sensitive member servers and administrative workstations.
- The IAO should delegate control of OUs to groups instead of individual accounts. This implementation of RBAC simplifies security administration.

As noted, delegating control over OUs creates limited-scope administrators. Ensuring that this privilege is controlled requires the maintenance of baseline documentation and periodic reviews. If an unauthorized user is able to gain control over an OU, that user could change the security policies applicable to the OU. This could result in weak security policies that would increase the risk of a successful attack against the affected resources.

- *(DS10.0260: CAT II) The IAM will limit the number of users and document those users who have been delegated AD object ownership or update permissions but are not members of Windows built-in administrative groups.*

Another potential use of delegation within AD is the case of GPOs. Members of the Group Policy Creator Owners group can create GPO objects in the AD database. This privilege can also be delegated to other user accounts through the GPMC tool. As with delegation of OUs, delegation of GPO creation can be used to create an environment in which fewer users have privileges on a specific collection of objects.

While delegating object creation can be advantageous, there are also risks associated with it. An inexperienced or malicious user might use the capability in a way that causes a denial of service for the domain controller. This can happen when so many objects are added that the server runs out of space for the AD database. In this case, it may become impossible to make additions or changes to the database.

A new feature was added to Windows Server 2003 as a mitigating control to this type of attack. This feature is the ownership quota for AD objects. It limits the number of objects that a user, group, computer, or service account is allowed to own in a specific AD database partition. The quota can be implemented two ways: as a default limit that applies to the partition for any account without a specific quota, or as the limit for the partition for a specific account.

If an unauthorized user is able to add an unlimited number of AD objects such as GPOs, a denial of service could occur for the affected domain controller. Logically, users who are unable to gain membership in an administrative group are more likely to attempt this kind of attack. Also, the domain partition is most likely to be vulnerable because delegation capabilities are commonly granted on objects in that partition of the AD database.

Microsoft does not currently offer guidance for a specific quota number because implementation environments may differ substantially. Also, because the number of tombstones (deleted objects) is a factor in the object ownership calculation, the volatility of a specific environment can affect the quota. Therefore the following guidance is recommended:

- In installations where the ability to create AD objects is delegated to users and groups who are not members of Windows built-in administrative groups, the SA should evaluate ownership patterns and define an appropriate quota for AD domain partition objects.

### **C.3.3.6. Functional Configuration - AD**

Section 3.3.7, Functional Configuration, states general requirements for replication scheduling, database integrity, account lockout, and OS services. The following information identifies any control settings needed for AD environments to comply with those requirements.

The concept of an AD site was briefly discussed in Section C.2.2, Forest and Domain Architecture. Defining AD sites has an impact on AD data replication. Although most replication within an AD domain occurs automatically and as required, the existence of multiple sites introduces an additional consideration. After defining AD sites, site links must be created to define the connections between them. Site links have schedule and replication interval properties that impact the time and frequency with which replication between sites is attempted.

If replication between domain controllers at different sites does not occur on a timely basis, information updated on the domain controllers in one site is not available to the domain controllers in the other site. Although the impact varies according to the data, this is generally undesirable if it continues for more than one day.

To comply with the replication requirement (DS00.3230) in section 3.3.7 for a forest implementation that includes multiple AD sites, the schedule and replication interval properties of each site link must be collectively configured to allow replication at least daily.

In order for AD processes to execute on a domain controller, there are certain Windows services that must be running. Without these services, group policy is not available when clients log on to the domain, the server cannot replicate with other servers, and Kerberos authorization processes may fail to function or function improperly.

Service-dependent AD processes can be disrupted by changing the startup type specification for the required Windows services. If the startup type is changed maliciously or unintentionally from “automatic” to “manual” or “disabled”, AD processes on the affected domain controller may be unavailable.

To comply with the requirement (DS00.3260) in section 3.3.7 for automated startup of OS services, the startup type for the following Windows services must be set to automatic on all domain controllers: Distributed File System, DNS Client, File Replication Service, Intersite Messaging, Kerberos Key Distribution Center, and Windows Time.

### C.3.4 Physical and Environmental

Section 3.5, Physical and Environmental, discusses the need for special physical access restrictions for directory servers that are critical parts of the infrastructure used for identification and authentication. The following information identifies the requirement for special physical access control for specific AD servers.

Domain controllers that hold the FSMO roles in the AD forest root domain are particularly important to the availability and integrity of resources on Windows servers throughout the forest. Data such as that in the AD schema is updated through a specific FSMO server and presents an attractive target for attacks employing poisoning or corruption strategies designed to impact an entire forest.

If an unauthorized user is able to obtain physical access to the servers holding the FSMO roles in the forest root domain, that user could more easily disrupt the operation of the entire forest. The impact of this disruption can vary widely from a short-term outage of a particular function to destruction that could require complete restoration of the domain and delays to inter-forest resource access. Because of this potential significant impact, access to these specific servers must be more restricted than for most.

- *(DS10.0310: CAT II) The IAO will ensure physical access to the forest root domain controllers that hold FSMO roles is restricted to specifically authorized personnel.*

### C.3.5 Continuity

Section 3.6, Continuity, identifies general requirements to support backup and restore processes for directory services. The following information provides guidance on meeting the general requirements, an additional requirement, and recommendations that are unique to AD for these processes.

In Section C.2.2, Forest and Domain Architecture, it is noted that AD data is backed up as part of an operation that backs up the Windows System State data. Among other items, a System State data backup on a domain controller includes the AD database and the GPT (SYSVOL) data that are necessary to restore the AD directory data.

To comply with the directory database backup requirement (DS00.0160) in section 3.6, it is necessary to perform a Windows System State data backup on each domain controller.

The DSRM system state used for AD restore operations was mentioned in Section C.2.2, Forest and Domain Architecture. When a Windows server is promoted to be a domain controller, a password must be selected for use when that domain controller is booted into DSRM. If that password is not available when needed, it will not be possible to boot the server and restore the AD database from the backup.

- *(DS10.0320: CAT II) The IAO will ensure the password defined for use in DSRM for each domain controller is stored in a locked, fire-rated container or is subject to other appropriate protections from loss.*

To correctly restore an AD domain controller, it is necessary to know what position it occupies within the AD implementation architecture. This refers to the domain controller's place within the AD forest as it relates to the domain and tree in which it resides. There may be other domain controllers within the domain as well as parent or child domain controllers that have to be factored into a recovery process.

To comply with the directory architecture details requirement (DS00.6120) in section 3.6, the following information is needed:

- Text or preferably a graphic representation of the AD forest that indicates the hierarchical relationship of the domains and trees
- A list of the domain controllers that hold FSMO and GC roles
- Indications of which domain controllers are configured to support SSL
- The suffixes (naming contexts) covered by the domains.

The quantity and placement of AD domain controllers can have a significant impact on recovery procedures. The general requirement (DS00.6140) in Section 3.6, Continuity, specifies the need for a redundant server, but the special considerations for AD are addressed here.

The large number of individual needs and environmental constraints within an organization the size of DoD makes it impossible in a practical sense to define a comprehensive list of strict AD server placement requirements. A configuration appropriate to a large Component presence could be unreasonable in cost and complexity terms for another. As a result, the following guidance is written as recommendations for the Components to interpret for their particular environment.

The following guidance should be applied when determining when to deploy servers and how to assign AD FSMO and GC roles:

- The IAO should ensure the Domain Naming and Schema Master roles are assigned to domain controllers that are not holding any other FSMO roles.
- The IAO should ensure the Infrastructure Master role is assigned to a domain controller that is not a GC server unless there is only one domain in the forest or all domain controllers are GC servers.
- The IAO should consider assigning the PDC Emulator role to a domain controller holding no other FSMO roles.
- The IAO should ensure at least one domain controller per AD site is a GC server. This should be increased to at least two domain controllers per AD site where GC-dependent server applications such as MS Exchange are being used.
- The IAO should ensure at least one domain controller per domain, that holds few or no FSMO roles, is designated as a standby operations master to assume the other roles.



## **APPENDIX D. ACTIVE DIRECTORY APPLICATION MODE SPECIFIC ELEMENTS AND REQUIREMENTS**

### **D.1. Introduction**

This appendix is reserved to address the technology-specific background and security requirements for ADAM. At this time the content is minimal. There is some guidance on meeting general directory service requirements and one requirement that is unique to ADAM. Additional guidance may be provided in future versions of this document.

ADAM provides a generic, LDAP-accessible directory service that runs in a user security context rather than an OS security context as does AD. Multiple instances of ADAM can run on a single host. ADAM was initially available as a downloadable feature that ran on Windows Server 2003 or Windows XP. ADAM now ships with Windows Server 2003 R2 as an optional component.

ADAM has been used within DoD as an application directory. As a native Windows directory service based on AD and potentially linked to AD data, ADAM may be easier to deploy by infrastructure groups already familiar with Microsoft directory architecture.

In 2005 the DoD Active Directory Interoperability Working Group (DADIWG) tested and documented scenarios demonstrating the use of ADAM in an AD synchronization solution. In this instance, ADAM was deployed as a border directory between two organizations. Using the MIIS product, one organization extracted directory data from their AD and populated an ADAM instance. The second organization used MIIS to review or filter that ADAM data and imported the result into their AD.

### **D.2. Active Directory Application Mode Security Background**

Because of the relationship between ADAM and AD, there are many common security issues and some that are unique. The following security considerations apply to the use of ADAM:

- The physical file components of ADAM are similar to those of AD. There are database, transaction log, and work files. OS-based file permissions are configured to protect these files.
- Objects in the ADAM database require appropriate access control.
- Although some common tools can be used, there are some unique administration tools for ADAM. This includes the ADAM ADSI Edit MMC snap-in (ADAM-adsiedit.msc) and versions of the Dsdbutil, Dsdiag, and Dsmgmt programs. Generally these tools can also be used to administer AD.
- Running on Windows XP, ADAM does not support auditing or password policies and the ADAM service account must be a member of the local Administrators group.
- User accounts can be defined within ADAM; they are one type of ADAM security principal. Windows security principals can also be defined as ADAM security principals. Groups defined within ADAM can contain ADAM security principals and Windows security principals.

- ADAM supports a special object that is identified as a proxy object. An ADAM proxy object represents an AD security principal and is used in bind redirection. The proxy object contains the SID of the AD security principal. When a bind to ADAM is performed using the proxy object, ADAM sends the SID and password to AD for authentication. An ADAM proxy object can be used to enable access to data stored in AD as well as ADAM.

### D.3. Technology-Specific Security Requirements

This section is reserved to supplement the information in Section 3, Directory Service Security Requirements. At this time only minimal coverage is provided for issues unique to ADAM.

#### D.3.1. Enclave and Computing Environment

Section 3.3, Enclave and Computing Environment, discusses multiple requirement areas related to the IA Controls in the Enclave and Computing Environment subject area. An additional requirement and specific guidance on meeting some of the general requirements for ADAM are discussed here.

##### D.3.1.1. Data Access Control - ADAM Files

Section 3.3.3, Data Access Control - Files, specifies the general requirement for access control for directory service data files. This section identifies the unique file system objects for ADAM.

ADAM utilizes a file structure similar to AD. This includes an ADAM database, log, and work files for each instance on a server. To comply with the requirement (DS00.0120) in section 3.3.3, it is necessary to configure the permissions included here.

Component	Object	Name	Type	Access
Database	...\adamntds.dit	Administrators SYSTEM [ADAM service account]	Allow Allow Allow	Full Control Full Control Full Control
Log files	...\edb*.log	Administrators SYSTEM [ADAM service account]	Allow Allow Allow	Full Control Full Control Full Control
Work files	...\temp.edb ...\edb.chk	Administrators SYSTEM [ADAM service account]	Allow Allow Allow	Full Control Full Control Full Control

It should be noted that the full path to these files has to be determined from the Windows registry entries for each ADAM instance. This is due to the fact that each instance has separate data files, and that other products utilizing Microsoft's Extensible Storage Engine create files with some of the same names.

### **D.3.1.2. Data Change Auditing - ADAM**

A special requirement related to audit capability applies to the use of ADAM. Currently Microsoft documentation states that auditing is not supported when ADAM runs on Windows XP Professional. The lack of audit capability makes the use of ADAM on this OS unacceptable for production applications.

- *(DS15.0100: CAT I) The IAO will ensure instances of ADAM used to process production data are not hosted on Windows XP-based computers.*

### **D.3.1.3. Group Membership and Limiting Privileges - ADAM**

Section 3.3.6, Group Membership and Limiting Privileges, specifies the general requirement (DS00.3210) to configure accounts used by directory service processes with the least privileges technically feasible. The following information addresses this requirement for ADAM.

Microsoft documentation indicates that ADAM instances can run using the Windows built-in Network Service account or another account with limited privileges. On this basis, ADAM instances must not run using an account that is a member of a local or domain Administrators group.

### **D.3.2. Continuity**

Section 3.6, Continuity, identifies general requirements to support backup and restore processes for directory services. The following information provides guidance on meeting two of the general requirements for ADAM.

Unlike AD, ADAM does not require certain OS-related data to be captured during a backup of the directory. To comply with the directory database backup requirement (DS00.0160) in section 3.6, it is only necessary to perform a data backup that includes all the files in the data directory of each ADAM instance. Microsoft states that the backup can be performed using the Windows Backup utility or any Windows Logo Program, third-party backup utility.

Groups of ADAM instances can participate in a defined “configuration set” in which a common schema, configuration, and any specified application directory partitions are replicated. ADAM supports the multi-master replication model among ADAM instances.

Like other directory services, the implemented directory architecture can impact the procedures used when it is necessary to restore a directory service. For this reason it is necessary to maintain documentation of the architecture. To comply with the directory architecture details requirement (DS00.6120) in section 3.6, the following information is needed:

- Text or preferably a graphic representation of the ADAM configuration set, including details such as the number of servers and replication flow.

This page is intentionally left blank.

## **APPENDIX E. RED HAT DIRECTORY SERVER SPECIFIC ELEMENTS AND REQUIREMENTS**

### **E.1. Introduction**

This appendix addresses the technology-specific background and security requirements for RHDS. The discussion of security background elements provides high-level technical information about aspects of RHDS that have security considerations. The security requirements section provides guidance on meeting the general directory service requirements for an RHDS installation, and additional requirements and recommendations that are unique to RHDS.

RHDS is a general purpose directory services solution. That is, it is not tightly integrated with a NOS in the way of directory service products such as AD. This is a significant distinction because it means that a given implementation of RHDS may, or may not, be tightly linked to OS or platform access security. When RHDS is used exclusively as a stand-alone directory or integrated with an application, a directory compromise may have less impact on an organization. On the other hand, if an RHDS server is used as a repository for identification data used by the Pluggable Authentication Module (PAM) implementation on one or more servers, a directory compromise could have significant and widespread impact.

The current version of RHDS is a descendent of the Netscape Directory Server product. Development at Netscape spanned several years and included collaboration with Sun in the iPlanet alliance. Although there was once a common code base, the end of the alliance resulted in different products with a common origin. This is important to remember when comparing security issues or configuration settings between the Red Hat and Sun products.

The importance of RHDS to DoD is reflected by the DoD-wide Red Hat Security Solutions Enterprise Software License. This agreement allows for the use of some Red Hat products, including RHDS, across the DoD, Intelligence agencies, and the CIA. Information about the agreement is available at <http://iase.disa.mil/netlic.html>.

On the subject of compliance with standards, Red Hat states that RHDS provides an LDAP server that is “compliant with the LDAPv3 Internet standards”. Because of the relatively recent changes in the standards and scope of the subjects in the newer proposed standards, it is assumed that this is a general statement that does not apply to every aspect of the standards or RHDS. In any case, RHDS does address the key security issues related to mechanisms for authentication and data integrity and confidentiality.

### **E.2. Red Hat Directory Server Security Background**

The security elements in RHDS are generally consistent with those discussed in Section 2.2, Common Directory Server Elements. The information here describes security aspects of the specific product implementation of the general elements discussed earlier.

Because RHDS is not tightly integrated with a NOS, there is no discussion of OS security issues in this appendix. The information presented here is based on the assumption that the OS guidance provided in other DoD-approved security configuration guidelines is implemented. For RHDS implementations on UNIX/Linux platforms, this includes specifically the UNIX STIG.

This appendix is not intended as a comprehensive source of information on RHDS. Only information that is thought to be minimally necessary to identify RHDS security elements is included. The Red Hat documents must be considered the primary source of product information. It is also noted that this document is not intended as a tutorial. It is assumed that persons attempting to comply with the stated requirements already have a sound understanding of the platform OS and RHDS.

In this appendix references are made to Red Hat's *Administrator's Guide, Red Hat Directory Server*. It provides critical information necessary to understand and administer RHDS. That document and the following documents provided primary input and background to this document:

- Red Hat's *Configuration, Command, and File Reference, Red Hat Directory Server*
- Red Hat's *Deployment Guide, Red Hat Directory Server*
- Red Hat's *Gateway Customization Guide, Red Hat Directory Server*

These and other associated documents are listed in Appendix A, Related Publications.

With the preceding information in mind, the next subsections provide brief descriptions of RHDS elements for which security considerations exist. Following that, the RHDS-specific security requirements are described in Section E.3, Technology-Specific Security Requirements.

### **E.2.1. Account and Group Considerations**

As a directory server, RHDS commonly stores user account information. As part of the product architecture, there are a number of accounts and groups that are defined in every installation. It is important to understand these accounts and groups because of the privileges that they have as a result of the server software, configuration controls, and default permissions.

Password control, group update integrity, and cross-directory authentication are also important issues for RHDS accounts. Without adequate consideration of these issues, accounts defined to an RHDS directory might not be used in accordance with the intended security policy.

#### **E.2.1.1. Accounts**

There are three accounts that are defined by default in an RHDS installation. A fourth type of account may be defined when replication is configured.

The Directory Manager account is created during RHDS installation and is the most highly privileged RHDS account. The Directory Manager has unlimited access to directory data; access control rules are not checked for access performed using this account. The default DN is "cn=Directory Manager" and the account and its password are attributes of the configuration subtree in the directory. Because of its access capability, this account is likely to be needed for some recovery scenarios.

The configuration directory administrator ID is created during RHDS installation and is automatically a member of the Configuration Administrators group. The directory administrator account receives access to all directory data. The default ID (uid) for the account is "admin".

The Administration Server user ID is created during RHDS installation and has the authority to administer all the servers defined in the local server root. This account is normally identical to the configuration directory administrator account and the password is kept in one of the Administration Server configuration files.

When the directory implementation includes servers that are updated through replication, there may be a distinct Replication Manager account (also known as the supplier bind DN) defined on each server that receives updates. This account is defined on these servers, known as hub or consumer servers, in a part of the database that is not replicated.

The credentials for the Replication Manager account are specified in the replication agreement settings on the supplier server and used by that server to access the hub or consumer server. The replication settings on the hub or consumer server identify the Replication Manager account. Access by the Replication Manager account is not subject to access control checking on the hub or consumer server so this account is considered a highly privileged user. The Replication Manager account is subject to password expiration that is defined in the global password policy, so action is necessary to ensure replication does not fail due to invalid credentials.

Instead of a distinct Replication Manager account, it is possible to specify that a supplier server identifies itself using its PKI certificate. This has three advantages. It eliminates the need to create a Replication Manager account on the hub or consumer server. It reduces the risk that a highly privileged account can be compromised. And it reduces the administrative burden of updating passwords that are stored in multiple locations. However, it does mean that the servers participating in replication must be configured to support SSL.

Recommendations for using account names other than the defaults are addressed in Section E.3.2, Identification and Authentication.

#### **E.2.1.2. Groups**

There are a few privileged groups that are defined by default in each RHDS installation: Configuration Administrators, Directory Administrators, and Server Instance Entry (SIE).

The Configuration Administrators group is defined in the NetscapeRoot subtree of the directory. The default access permissions allow this group to update the configuration, schema, NetscapeRoot, and user root subtrees in the directory.

The Directory Administrators group is defined in the user root subtree(s) of the directory. The default access permissions allow this group to update the configuration, schema, and user root subtrees in the directory.

The SIE group is defined in the NetscapeRoot subtree of the directory. The default access permissions allow the members of this group to update the configuration, schema, and user root subtrees in the directory. The members of this group are the RHDS instances (directory servers) that are created within the server root file system.

It is necessary to monitor membership in privileged groups to ensure intended directory access is maintained. Requirements for this are addressed in Section E.3.3.5, Group Membership and Limiting Privileges – RHDS.

In most cases local administrators define additional account groupings within a directory. RHDS allows different types to support flexibility and legacy compatibility. Static and dynamic groups are provided for compatibility with older versions of RHDS. Managed and filtered roles are now the preferred account grouping mechanism. The groupings are described as follows:

- Static groups – A static group is a collection of accounts based on specifically identified members. Each static group directory entry contains a uniqueMember attribute for each account in the group.
- Dynamic groups – A dynamic group is a collection of accounts based on an attribute filter that is applied to the directory at the time of reference.
- Managed roles – A managed role is a collection of accounts based on specifically identified members. Each account contains an nsRoleDN attribute for each role it is assigned.
- Filtered roles – A filtered role is a collection of accounts based on an attribute filter that is applied to the directory at the time of reference.

Please note that it is possible to create nested groups and roles; that is a group within another group or role within another role. However, there is no unique security implication in doing so.

Beyond the obvious consideration of controlling membership in groups or roles, it is important to control a user's access to update the attributes used to assign membership. Requirements for this are addressed in Section E.3.3.5, Group Membership and Limiting Privileges – RHDS.

A default installation of RHDS may also include non-privileged groups that are provided as possible bases for local customization. These accounts are addressed by guidance in Section E.3.2, Identification and Authentication.

### **E.2.1.3. Account and Group-Related Controls**

Controls are available within RHDS to improve the integrity of account and group data. By enforcing stronger password practices, automatically updating group membership data, and supporting limited cross-directory authentication, RHDS has features that can support a higher level of security for the accounts defined in each directory.

RHDS supports the configuration of password controls at three levels. The global password policy can be implemented through values assigned to several core server configuration (cn=config) attributes. Password age, length, encryption algorithm, and composition (to a limited extent) can be controlled.



It is also possible to define a password policy entry and apply it to a subtree or user level. The same attributes available at the global level can be applied to a subtree or user. This type of control may be appropriate for application (that is non-user) accounts that are defined within the directory. For those accounts, it may be desirable to specify a longer password age than is applied in the global password policy.

Requirements for password policy controls are addressed in Section E.3.2, Identification and Authentication.

In the previous discussion of static groups and managed roles, the relationship between these grouping mechanisms and accounts was discussed. Because account entries and group or role entries are related, it is critical that updates to one are properly reflected in the other. In conventional database terms this is referred to as a referential integrity issue. RHDS provides the Referential Integrity Plug-in to address this.

Additional information and implementation guidance for the Referential Integrity Plug-in is located in Section E.3.3.6, Functional Configuration – RHDS. It is important to review the implementation considerations discussed there to avoid performance problems and replication issues.

Cross-directory authentication can be a powerful tool to allow users defined in one directory to access resources defined in, and controlled under, another directory. It might also be used to support a single sign-on strategy so that fewer accounts need to be defined. The Pass-through Authentication (PTA) mechanism, implemented through the PTA Plug-in in RHDS, provides support for cross-directory authentication.

A simple implementation of PTA is automatically set up by RHDS when a configuration directory and user directory are installed in separate directory server instances. This allows the configuration directory administrator to update the user directory. The following briefly describes access authorization in this scenario:

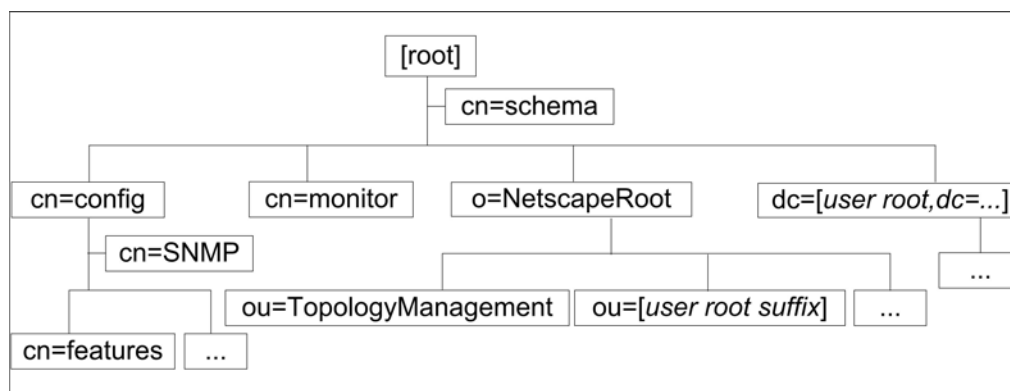
- The configuration directory administrator supplies his credentials during a bind to the user directory server.
- The PTA configuration on the user directory server directs it to transmit the supplied credentials to the configuration directory server where the account is defined.
- The configuration directory server authenticates the configuration directory administrator's credentials and responds to the user directory server.
- Following a successful reply from the configuration directory server, the user directory server allows the configuration directory administrator to bind.

Beyond this simple implementation, PTA could be configured to allow cross-directory authentication between multiple user directories. While this can be a valuable directory implementation strategy, safeguards are necessary to ensure it is used to enable access that is consistent with intended security policy. Section E.3.3.1, Cross-Directory Authentication (Pass-through Authentication), specifies requirements for the implementation of PTA.

## E.2.2. Object Access Control

Object access control is clearly a critical security requirement for any directory. Without it, there can be no expectation that the confidentiality, integrity, or availability of directory data can be maintained. Because directory data is often used to make access control decisions, a compromise of directory data is likely to lead to compromises of other data. The primary way in which RHDS implements object access control is through the ACI mechanism.

Before discussing the RHDS ACI implementation, it is important to describe the high-level, logical view of an RHDS directory tree. This is important so that the placement and scope of ACIs can be understood. The following figure is an abbreviated high-level view of an RHDS directory tree.



**Figure E-1. RHDS Tree**

The highlights of this structure are:

- The configuration subtree (cn=config) contains the core directory server configuration attributes. Many of the attributes in this subtree have an impact on the security of directory server instances. The data in the subtree is read from the `dse.ldif` file when a directory server is started, and written to the `dse.ldif` file when the server is stopped.
- The monitor subtree (cn=monitor) contains server performance data.
- The NetscapeRoot subtree (o=NetscapeRoot) contains configuration attributes for the Administration Server and the entries for the configuration directory administrator account and the Configuration Administrators group.
- The user root subtree(s) contain the primary directory data, including entries for the Directory Administrators group.
- The schema object (cn=schema) contains the schema definitions currently used by the directory server. Data in this object is built from the LDIF files in the `serverRoot/slapped-serverID/config/schema` file system directory.

RHDS provides the capability to control object access at a detailed level throughout the directory tree. Permissions can be defined for the entire directory, a subtree, an entry, and an entry attribute. Permissions are defined at a specific level in the tree and are inherited downward from that entry. Permissions can be made conditional on an attribute of the account being used to perform the access. This level of control is applied through the creation of ACIs that reside in the directory.

The following notes apply to the use of access control in RHDS:

- The Directory Manager account is not subject to access control.
- When replication is configured, the Replication Manager account (also known as the supplier bind DN) is not subject to access control on the consumer server.
- When an ACI grants proxy rights to an account, that account can be used to impersonate any account for access to the data covered by the ACI.
- Although the RHDS documentation indicates a default that “all users are denied access rights of any kind”, a broad range of privileged and anonymous access is configured by the default ACIs created during server installation.

The subject of ACIs is a critical, but complex part of RHDS security administration. The following discussion only provides very high-level information. Please refer to the chapter on access control in Red Hat’s *Administrator’s Guide, Red Hat Directory Server* for details on creating and implementing ACIs.

In most cases ACIs can be created through the administrative console or through the application of LDIF files. However, the console interface does not support all the available ACI options. In either case, each ACI includes the following parts:

- The ACI target specifies the portion of the directory to which the ACI applies. In the most general terms this usually means an entry, all the entries in a subtree, an attribute, or some combination. Multiple attributes can be specified, but only one entry. Instead of a specific entry, a search filter can be used to indicate the attributes of multiple entries are controlled.
- The ACI permission specifies the type of access that is being allowed or denied. The individual permissions are: read, write, search, compare, self-write, add, delete, and proxy. A permission of “all” indicates that every permission except proxy is included.
- The ACI bind rule specifies the DN or network location to which the ACI applies. Wildcards can be used in some cases, most notably for network specifications. A search filter can be used to indicate that all accounts that fit the filter criteria are subject to the ACI.

Certain options of the ACI bind rule have very significant security implications:

- DN-level access can be specified at a user basis (userdn), group basis (groupdn), role basis (roledn), or attribute (userattr) basis.
- User-based bind rules can include the following keywords:
  - The “all” keyword indicates the rule applies to any user who has successfully authenticated to the directory.
  - The “anyone” keyword indicates the rule applies to any user, whether authenticated or not. This keyword is used to enable anonymous access.

- The “self” keyword indicates the rule applies if the bind DN matches the DN of the target entry.
- The “parent” keyword indicates the rule applies if bind DN is the parent of the target entry.
- Time of day or day of week can be used to specify that the rule applies to specific periods of time.
- The method used for authentication, specified as “none”, “simple”, “SSL”, or “SASL”, can be used to specify that the rule applies to bind connections using the authentication mechanism identified. Used in a certain manner, the “none” specification might be used to enable anonymous access.

The following notes apply to the placement of ACIs within the RHDS directory tree and the resulting evaluation:

- ACIs placed at the root DSE entry apply only to that entry.
- ACIs that target an entry that is a branch point in the directory apply to the branch point as well as all of its child entries. In other words, permissions are inherited downward.
- When access to an entry is being evaluated, the proximity of the entry to an ACI is not a factor. Instead the ACIs for the entry and all the parent entries back up to the top level are evaluated.
- If conflicting permissions exist, the Precedence Rule is applied. The Precedence Rule states that the permission that denies access always takes precedence over the permission that allows access.
- ACIs are stored as attributes of entries, placed at the point of an entry within the directory tree. As a result, access control is impacted by replication.

The considerations for the impact of replication on ACIs are:

- ACIs are replicated like other attributes. This means that security policy is enforced across replicated directories on multiple servers.
- Inappropriate configuration of the replication schedule impacts the timeliness of changes to ACIs.
- The use of fractional replication impacts access control when the replicated portions do not include the higher levels at which ACIs are defined or do not include other attributes which are used in the ACI definition.

It is possible to examine the access control that results from the defined ACIs through the “get effective rights” LDAP control that is implemented in RHDS. This control is invoked through the console or through the “-J” option on the ldapsearch command. The target directory entry and the proposed accessing account are input. The access rights to the entry and to its attributes are returned. The *Administrator’s Guide, Red Hat Directory Server* has a detailed description of the get effective rights control.

Section E.3.3.3, Data Access Control – RHDS Directory Database Objects, specifies basic requirements for controlling access to directory data and some specific requirements for the use of ACI capabilities.

### **E.2.3. Other Implementation Features and Details**

RHDS offers advanced features that make it a robust directory server. Some of the features provide optional capabilities that are not implemented in every directory environment. On the other hand, there are some product implementation details that are inherent to every RHDS installation. In any case, it is important to understand that there are security impacts that must be considered when the product implementation details and features are configured.

The following subsections provide a high-level discussion of some of the RHDS implementation features and details that have security concerns. This includes replication, administrative interfaces, plug-ins, gateways, chaining, and ports and protocols. This information is intended to provide background. The vendor manuals must be consulted for complete information.

#### **E.2.3.1. Replication**

Replication is the process whereby directory data is copied from one server to another. Replication is used to enhance the availability of directory data and the performance of directory queries. RHDS provides sophisticated capabilities that include multi-master, cascading, and fractional replication.

There are several terms that are used in describing replication:

- Supplier servers hold a read-write replica of directory data and push the data to consumer servers.
- Consumer servers hold a read-only replica of directory data and are updated by supplier servers.
- Hub servers are used in cascading replication configurations. A hub server acts as a consumer server, holding a read-only replica that it receives from a supplier server, and as a supplier server, providing updates to consumer servers.
- Multi-master replication implies that multiple servers allow updates to the same directory data and a resolution process exists for managing conflicts.
- Cascading replication provides a more efficient way of updating multiple consumer servers by allowing hub servers to act as distribution points.

Although replication items are configured on each server that participates, the “replication agreement” entry that is configured on a supplier server holds many of the critical configuration attributes. A replication agreement identifies the directory database being replicated, entry attributes to be included, the consumer server, available time periods, the credentials to use to connect to the consumer, and connection security.

There are controls that can be configured to ensure replication connections are secure:

- It is possible to configure mutual authentication between the consumer and supplier servers through the use of SSL.
- The credentials of the supplier server can be provided to the consumer server through the Replication Manager account (also known as the Supplier Bind DN) or through the supplier server’s PKI certificate.
- The confidentiality and integrity of directory data in transport can be protected by specifying the use of SSL for the connection.

Requirements and implementation guidance for replication are included in Section E.3.2, Identification and Authentication, Section E.3.3.6, Functional Configuration – RHDS, and Section E.3.3.7, Data Transmission Confidentiality and Integrity – RHDS.

### **E.2.3.2. Administrative Interfaces**

There are three principal interfaces for administrative access to instances of RHDS. In the context of security, administration refers to accessing configuration controls and data that require special privileges and that can affect the operation of the server and the security of the data. The three interfaces are: editing configuration files, using command line utilities such as `ldapsearch` and `ldapmodify`, and using the Red Hat Console.

A text editor can be used on the server host to directly read or update files that contain configuration settings for a directory server and an administration server. The principal file for a directory server is `dse.ldif`, but there are other files such as `99user.ldif` (or other schema update files) that can be manually edited. As with other configuration files in UNIX systems, the chief way in which access to these files is protected is file permissions.

The use of command line utilities for administration refers commonly to the `ldapsearch` and `ldapmodify` programs that are provided with RHDS. As the names imply, `ldapsearch` is used for directory queries and `ldapmodify` is used for updates. These programs can use the LDAP or LDAPS protocol over ports 389 or 636 to communicate with the directory server. However, because RHDS is LDAP-compatible, there are other versions of these tools as well as different tools that could be used to access RHDS. In the absence of GUI-based tools, non-privileged directory server clients as well as administrators may validly use the `ldapsearch` utility for directory searches.

It is not generally practical to implement technical controls on the command line utilities because they often reside on client machines and may be used by clients performing legitimate directory queries. The chief ways in which access from command line tools is controlled is through administrative policies, firewalls that limit port access, and limits on authentication methods supported by the directory server.

The Red Hat Console provides the commonly used graphical interface to perform RHDS administrative tasks. [Red Hat Console is actually the front-end management application used by multiple Red Hat products.] The console is automatically installed, along with the Administration Server software, when the first RHDS instance is installed on a host. The console can also be installed on other hosts. The console itself is a Java-based client application that connects to both an administration server and a directory server instance.

The console connects to an administration server instance using the HTTP or HTTPS protocol, according to the server's configuration. The administration server user ID and its password are usually used to access the server. Using the directory configuration data held by the administration server, the console can also connect to the directory server to perform administrative tasks.

Administrative traffic that flows from the console includes a password for the initial connection to the administration server and may subsequently include other sensitive data (such as account passwords) over the connection to the directory server. For this reason it is essential that encryption is employed when these connections traverse a network.

Implementation guidance for network traffic encryption is included in Section E.3.3.7, Data Transmission Confidentiality and Integrity – RHDS.

### **E.2.3.3. Plug-ins**

RHDS supports plug-ins to provide enhancements or extensions to the functions of the directory server. Almost three dozen plug-ins are shipped with RHDS and normal operation of a directory server requires that many of these be enabled at all times.

The following plug-ins have special relevance to security:

- The PTA Plug-in enables pass-through authentication to be performed. When this involves cross-directory authentication (beyond configuration and user directory separation), there may be security issues depending on the classification levels of the directories or a connection with a non-DoD organization. Also, protection of authentication data in transit has to be considered.
- The Referential Integrity plug-in is involved in maintaining database integrity when updates of group or role memberships take place. When group or role-related updates are performed and this plug-in is not enabled and properly configured, group or role membership might be incorrectly indicated in the directory.

Requirements related to the use of PTA are included in Section E.3.3.1, Cross-Directory Authentication (Pass-through Authentication). Implementation guidance related to the Referential Integrity Plug-in is included in Section E.3.3.6, Functional Configuration – RHDS.

### **E.2.3.4. Gateways**

Red Hat defines an RHDS gateway as “an HTTP-to-LDAP client that lives on an HTTP server”. In other words, an RHDS gateway provides a method by which clients can access a directory server through a web server. This can hide the complexity associated with LDAP query syntax and allow the greater number of clients with web browser software to access directory data. Multiple gateway instances can be defined on a single web server.

During direct server installation, an instance of each of the following is also installed:

- The Directory Express gateway provides basic directory query capability. Configuration data for this gateway type is located in a pb.conf file.
- The Default Gateway provides search, create, and modify capabilities. Configuration data for this gateway type is located in a dsgw.conf file.

The standard installation of RHDS configures the administration server as the web server for the Default Gateway and Directory Express gateway. This configuration is a security concern because it merges administrative and user functions in a single server. It could also increase risk

to the host when the administrative server runs with superuser credentials (to support privileged port access) and a vulnerability in a gateway component is discovered and exploited.

A DSML Gateway is also included with RHDS. This gateway allows web services clients to access directory data using DSML instead of LDAP. Configuration data for this gateway type is located in a `dsmlgw.cfg` file. Although it needs to be activated, the standard DSML Gateway implementation is based in the administration server as the other gateways are. For the same reasons as with the other gateway types, this is a concern. The default configuration of the DSML Gateway specifies the use of anonymous access. This is not consistent with the desired security policy for non-public directory data.

A requirement to restrict gateway implementation and implementation guidance to protect gateway configuration files are included in Section E.3.1, Security Design and Configuration. A requirement to alter the default anonymous access configuration setting is included in Section E.3.3.3, Data Access Control – RHDS Directory Database Objects.

#### **E.2.3.5. Chaining**

Chaining is one type of directory knowledge reference supported by RHDS. When chaining is configured, a client's request for data not on the directory server is relayed by the original server to another server where the information resides. Cascading chaining is a logical extension. It allows intermediate servers between the original and final servers.

RHDS enables chaining through the definition of database links. A link specifies the target server where data resides and the credentials to be used to access that server. The account associated with the credentials is granted the proxy access right to the data. If this is a privileged account on the target server, access to the target server's data could be compromised.

When cascading chaining is configured, intermediate database links are defined. A special link attribute, `nsCheckLocalACI`, is available to ensure the proper credentials are being evaluated through the chain. Details on cascading chaining should be reviewed in the *Administrator's Guide, Red Hat Directory Server*.

Requirements for the use of accounts defined in database links and the setting of the `nsCheckLocalACI` attribute are included in Section E.3.3.3, Data Access Control – RHDS Directory Database Objects.

#### **E.2.3.6. Ports and Protocols**

Section 2.2.5, Network Ports and Protocols, discusses the standard ports and protocols for services that directory servers provide and consume. RHDS utilizes these as well as another that is associated with an administrative interface.

A directory server instance is a network service provider as follows:

- The RHDS directory server listens by default on port 389 for traffic formatted in LDAP.
- When configured for SSL, the RHDS directory server listens by default on port 636 for traffic formatted in LDAPS.



- RHDS allows the ports for LDAP and LDAPS traffic to be configured as any number between 1 and 65535.

There are security implications to the use of ports 389, 636, or any port below 1024. When any of these port numbers is used, the directory server has to be installed and executed using superuser credentials. If a directory server executes using superuser credentials, the associated administration server must also run as superuser if the administration server is going to be used to start the directory server.

At the time an administration server is installed, a port number is selected. This is the port on which the administration server listens for connections from the Red Hat console. There is no default number for this port. If the installer does not specify a number, a randomly generated number is used. The assignment of this port can be an issue when the traffic from the console to the administration server traverses a firewall or other protected network boundary.

A directory server instance is a network service consumer as follows:

- When replication is in use, a supplier directory server initiates a connection to a consumer directory server, specifying the port in the replication agreement (nsDS5ReplicaPort attribute). This usually means that the supplier connects to port 389 or 636 on the consumer server.

### **E.3. Technology-Specific Security Requirements**

This section supplements the information in Section 3, Directory Service Security Requirements. It provides guidance on meeting the general requirements, additional requirements, and recommendations that are unique to the security considerations for RHDS.

It is important to understand specific terminology used in this section:

- Directory database - RHDS data is stored in a series of files that are collectively referred to as the directory database. This is not a database in the sense of a general-purpose data store, but it does provide the data repository for RHDS and so is referred to by the term directory database in this document.
- Subtree - A subtree is a portion of a directory tree. In RHDS, configuration information is stored in the NetscapeRoot and "config" subtrees. The user data for a directory is stored in one or more user root subtrees.

It should be noted that the requirements here are based on version 7.1 Service Pack (SP) 3 of RHDS. The requirements may be applicable to other versions, but no specific evaluation of other versions has been done. As with the implementation of all security configuration guidance, DoD Components should test configuration changes in a test environment before implementation in production to ensure their specific environment is not impacted in unintended ways.

### E.3.1. Security Design and Configuration

Section 3.1, Security Design and Configuration, discusses multiple requirement areas related to the IA Controls in the Security Design and Configuration subject area. There are unique considerations for RHDS in the requirement areas for configuration and implementation integrity and software integrity so details and guidance on meeting the general requirements for RHDS are discussed here. There is a special consideration for the security service partitioning area that dictates the need for an additional requirement for RHDS.

To comply with the requirement (DS00.1130) in section 3.1.2 to use FIPS-validated encryption, the following attribute values must be set when the server is configured to support SSL:

Entry DN	Attribute	Value
cn=encryption, cn=config	nsssl2	off
cn=encryption, cn=config	nsssl3	on
cn=encryption, cn=config	nsssl3ciphers	-rsa_null_md5,-rsa_rc4_128_md5, -rsa_rc4_40_md5,-rsa_rc2_40_md5, -rsa_des_sha,+rsa_fips_des_sha, +rsa_3des_sha,+rsa_fips_3des_sha, -fortezza,-fortezza_rc4_128_sha, -fortezza_null, -tls_rsa_export1024_with_rc4_56_sha, -tls_rsa_export1024_with_des_cbc_sha

**Table E-1. Cipher Attribute Values**

**NOTE:** Although the rsa\_3des\_sha algorithm is not FIPS-validated, it is being enabled because it is the strongest choice available that allows use of the Red Hat Console with SSL.

To comply with the requirement (DS00.1150) in section 3.1.4 to restrict access to directory software libraries and configuration files, the permissions in the following tables (or more restrictive ones) must be set.

**NOTES:** - The default value of *serverRoot* is /opt/redhat-ds or /opt/redhat-ds/servers.  
 - The value of *serverID* in *serverRoot/slapd-serverID* depends on the string specified during the server instance creation.  
 - The owner and group specifications annotated with an asterisk (\*) can be root or a dedicated directory server account\group. The directory server account cannot be the same account that is used to execute the server.

Component	Object	Owner	Group	Permissions
RHDS Server root	<i>serverRoot</i>	root	root	755
Server PKI data	<i>serverRoot/alias</i>	root *	root *	750
Executables	<i>serverRoot/bin/slapd/server</i>	root *	root *	700

Component	Object	Owner	Group	Permissions
Gateway configuration	<i>serverRoot</i> /clients/dsgw/context	root *	root *	755
Gateway configuration files	<i>serverRoot</i> /clients/dsgw/context/*.conf	root *	root *	640
Utility programs	<i>serverRoot</i> /shared/bin	root *	root *	755
Utility configuration	<i>serverRoot</i> /shared/config	root *	root *	755
Configuration	<i>serverRoot</i> /userdb	root *	root *	755

**Table E-2. Server Software Directories**

Component	Object	Owner	Group	Permissions
Admin Server root	<i>serverRoot</i> /admin-serv	root	root	755
Configuration	<i>serverRoot</i> /admin-serv/config	root *	root *	755
Configuration file	<i>serverRoot</i> /admin-serv/config/adm.conf	root *	root *	600
Administrator account file	<i>serverRoot</i> /admin-serv/config/admpw	root *	root *	600
Configuration file	<i>serverRoot</i> /admin-serv/config/local.conf	root *	root *	640
Configuration file	<i>serverRoot</i> /admin-serv/config/magnus.conf	root *	root *	600
Configuration file	<i>serverRoot</i> /admin-serv/config/server.xml	root *	root *	600

**Table E-3. Administration Server Software Directories and Files**

Component	Object	Owner	Group	Permissions
Configuration backup	<i>serverRoot</i> /slapd- <i>serverID</i> /confbak	root *	root *	750
Configuration	<i>serverRoot</i> /slapd- <i>serverID</i> /config	root *	root *	755
Schema configuration	<i>serverRoot</i> /slapd- <i>serverID</i> /config/schema	root *	root *	755
DSE files	<i>serverRoot</i> /slapd- <i>serverID</i> /config/dse*.ldif*	root *	root *	600

**Table E-4. Server Instance Software Directories and Files**

In section 3.1.5, Security Service Partitioning, there is a requirement to isolate directory servers from other applications including database and web servers. There is a special consideration for RHDS because of bundled software. The Default Gateway, Directory Express, and DSML gateway applications provided with the RHDS software are web server applications, installed by default to run using the administration server as host. In fact, the normal RHDS installation

process configures the Default Gateway and Directory Express applications to work without further administrator effort.

It is inconsistent with the isolation requirement to allow non-privileged users to access gateway applications on the same web server where administrative tasks are performed. It can expose the server to vulnerabilities in the applications and can be a significant risk when the server runs in its typical configuration with superuser credentials. To avoid these vulnerabilities and risks, non-privileged users cannot have access to gateway applications on the same host as the administration server.

- *(DS20.1100: CAT II) The IAO will ensure non-privileged users are not able to access any gateway applications (Default Gateway, Directory Express, and DSML) on the server on which the administration server is running.*

The Component site may choose the mechanism used to restrict access. The following are possible access controls:

- Removing the files under the *serverRoot/clients* directory on the administration web server or setting file permissions that prevent the web server from accessing the files (as long as the administration server is not executing under superuser credentials).
- Restricting network access to the port used for the administration web server such that access is available only to workstations used by privileged users (e.g., administrators).

### E.3.2. Identification and Authentication

Section 3.2, Identification and Authentication, describes general requirements for identification and authentication data protection for directory services. The following information provides details and guidance on meeting the general requirements for RHDS security for account and authentication data.

If directory server controls are being used to comply with the requirements (DS00.2110, DS00.2115) in section 3.2.1 for password complexity, expiration, and history, the following core server configuration (cn=config) attribute values must be set for the global password policy:

Attribute	Value
passwordCheckSyntax	on
passwordExp	on
passwordGraceLimit	0
passwordHistory	on
passwordInHistory	5
passwordMaxAge	5184000
passwordMinAge	86400
passwordMinLength	9
passwordStorageScheme	SSHA
nsslapd-rootpwstoragescheme	SSHA

**Table E-5. Password Policy Controls – Content, History, Age**

These values are intended to force passwords to have the following characteristics:

- Passwords are at least nine characters long.
- Password content does not contain user's name, ID, or any attribute value stored in the uid, cn, sn, givenName, ou, or mail attributes of the user's directory entry.
- Passwords expire after sixty days (5,184,000 seconds).
- A history of five previously used passwords is maintained and passwords can only be changed once per day.
- Passwords are encrypted using the strongest available algorithm.

These are minimum requirements that may be exceeded by local policy or when a more restrictive policy is in effect due to INFOCON status changes. Although this configuration does not enforce all the password content requirements for DoD, it is the strongest configuration currently available for RHDS.

RHDS supports granular password policy control through subtree or user level policies. There may be circumstances in which an organization needs to utilize this level of control for application accounts. Because of the vulnerability created by allowing weak passwords, any policy that is less restrictive than the global policy must be explicitly approved and documented.

- *(DS20.2100: CAT II) The IAM will ensure any subtree or user level password policies that are less restrictive than the global policy are approved in writing and the settings are listed in documentation for the directory service.*

There is a standard exception to one part of this policy for the password expiration for accounts used for directory replication. These accounts are generally known as Replication Manager or supplier bind DN accounts. The password expiration for those accounts can be configured to be one year or less.

The requirements (DS00.2120, DS00.2121) in section 3.2.1 to change factory set, default, or standard accounts and passwords are intended principally to ensure conventional, well-known passwords are not used with vendor-defined accounts. The installation and setup of RHDS does not employ pre-populated passwords, so the primary intent of the requirement is met. However, there are some account naming and group creation considerations.

A security benefit is gained by using non-standard names for certain, fixed RHDS accounts instead of the vendor defaults provided during installation. Choosing different names may reduce the threat from attacks based on the default names identified in the documentation. The following guidance addresses the fixed RHDS accounts:

- The SA should assign a CN other than "Directory Manager" for the root DN (also known as Manager DN) created during directory server setup.
- The SA should assign a uid other than "admin" for the configuration directory administrator ID created during directory server setup.
- If replication identity is based on an account rather than PKI certificate, the SA should assign a CN other than "Replication Manager" for the supplier bind DN used for replication.

The default installation of RHDS creates some groups that are not required for operation of the server. Although these groups do not automatically have any members, they do have some default access permissions and could provide a base on which an attack might be constructed.

Therefore the following default groups must be deleted to comply with the requirement (DS00.2121) in section 3.2.1: “Accounting Managers”, “HR Managers”, “QA Managers”, and “PD Managers”.

### **E.3.3. Enclave and Computing Environment**

Section 3.3, Enclave and Computing Environment, discusses multiple requirement areas related to the IA Controls in the Enclave and Computing Environment subject area. There are special considerations that dictate the need for additional requirements and implementation considerations for RHDS in the cross-directory authentication, data access control, data change auditing, group membership and limiting privilege, functional configuration, and data transmission confidentiality and integrity areas. The additional requirements and specific guidance on meeting the general requirements for RHDS are discussed here.

#### **E.3.3.1. Cross-Directory Authentication (Pass-through Authentication)**

Section 3.3.2, Architecture and Cross-Directory Authentication, provides a very general discussion of the impact of directory implementation architecture and cross-directory authentication. This section identifies the unique aspects of those issues for RHDS.

RHDS provides the PTA mechanism to allow one directory server to make calls to another directory server to perform bind authentication. The implementation of PTA to act as a single logical directory, such as with separate configuration and user directory servers, would not cause special concern. However, using PTA to bridge certain functional or organization boundaries can have serious implications.

It is a recommendation that cross-directory authentication should be disabled in most circumstances, but there are two specific instances in which it almost always must be disabled: directories for resources at different classification levels and directories for DoD and non-DOD organizations. Because there is not a known controlled interface to monitor directory traffic, directories at different classification levels must remain isolated. Because DoD and non-DOD organizations are likely to have different security controls applied and consequently different security postures, review and approval of those implementations is needed.

- *(DS20.3100: CAT I) The IAO will ensure PTA is not configured between directory servers at different classification levels.*
- *(DS20.3105: CAT I) The IAO will ensure PTA is not configured between DoD and non-DOD directory servers unless:*
  - *The network connections comply with all requirements for external connections defined in the Network Infrastructure STIG, including a Memorandum of Agreement (MOA) between the two parties*
  - *Explicit approval of the PTA configuration by the DAA is documented.*

Because the implementation of PTA involves transmitting authentication information between servers, it is necessary to use encryption to ensure the confidentiality and integrity of the data in transit over any network.

- (DS20.3110: CAT II) The IAO will ensure PTA configurations between physical hosts utilize LDAPS for the connection protocol or LDAP with host-to-host encryption.

### E.3.3.2. Data Access Control - RHDS Files

Section 3.3.3, Data Access Control - Files, describes the general requirement for access control for directory service data files. This section identifies the permissions for unique file system objects for RHDS.

The RHDS data files include many database, log, index, and lock files. To comply with the requirement (DS00.0120) in section 3.3.3, it is necessary to ensure the explicit permission specifications included here (or more restrictive) are set.

**NOTES:** - The default value of *serverRoot* is /opt/redhat-ds or /opt/redhat-ds/servers.  
 - The value of *serverID* in *serverRoot/slapd-serverID* depends on the string specified during the server instance creation.  
 - The owner and group specifications annotated with an asterisk (\*) can be root or a dedicated directory server account\group.

Component	Object	Owner	Group	Permissions
Directory Server root	<i>serverRoot/slapd-serverID/</i>	root *	root *	755
Database backup	<i>serverRoot/slapd-serverID/bak</i>	root *	root *	750
Database	<i>serverRoot/slapd-serverID/db</i>	root *	root *	750
Logs	<i>serverRoot/slapd-serverID/logs</i>	root *	root *	700
Referential Integrity log file (if RI-enabled)	<i>serverRoot/slapd-serverID/logs/referint</i>	root *	root *	700
Replication changelog (if multi-master replicating)	<i>serverRoot/slapd-serverID/changelogdb</i> [or Site-specified location]	root *	root *	700

**Table E-6. Server Instance Data Directories and Files**

### E.3.3.3. Data Access Control - RHDS Directory Database Objects

Section 3.3.4, Data Access Control - Directory Database Objects, describes the general requirement for access control for directory database objects. This section identifies specific database objects with required access settings and additional unique requirements for RHDS.

Section E.2.2, Object Access Control, briefly describes the ACI mechanism and how it applies to the directory tree structure in RHDS. To meet confidentiality and integrity requirements, ACIs defining access control are required at the root of the tree as well as at the subtree levels. The values created during the default RHDS install do not provide adequate access control for directory data at a non-public (classified or sensitive) confidentiality level. The following table summarizes the minimum required access control configuration for each level of the directory tree to meet the requirement (DS00.0130) in section 3.3.4.

<b>Object</b>	<b><u>Level</u></b> - <b>Account – Access</b> -- <b>Attribute Limits</b> (if any)
Root DSE (dn:)	Authenticated users - Read, Search, Compare
Config subtree (dn: cn=config)	<u>Base:</u> - Configuration Administrators group - All - administrator user account - All - Directory Administrators group - All - Server Instance Entry group - All <u>“cn=SNMP”:</u> - Authenticated users – Read, Search, Compare <u>“oid=2.16.840.1.113730.3.4.9,cn=features”:</u> [Entry for VLV index information] - Authenticated users – Read, Search, Compare, Proxy -- all attributes except “aci”
Monitor subtree (dn: cn=monitor)	<u>Base:</u> - Authenticated users – Read, Search, Compare -- all attributes except “aci” and “connection”
Schema attribute (dn: cn=schema)	<u>Base:</u> - Authenticated users – Read, Search, Compare -- all attributes except “aci” - Configuration Administrators group - All - administrator user account - All - Directory Administrators group - All - Server Instance Entry group - All



Object	<b>Level</b> <b>- Account – Access</b> <b>-- Attribute Limits (if any)</b>
NetscapeRoot subtree (dn: o=NetscapeRoot)	<u>Base:</u> - Authenticated users – Read, Search - Configuration Administrators group - All - group members - Read, Search, Compare [required to set server access by group] <u>“ou=TopologyManagement”:</u> - Authenticated users – Read, Search, Compare <u>“ou=[user root]”:</u> - Authenticated users – Read, Search <u>“ou=Global Preferences,ou=[user root]”:</u> - Authenticated users – Read, Search <u>“ou=UserPreferences,ou=[user root]”:</u> - Authenticated users - Read, Search, Compare <u>“cn=PublicViews,ou=4.0, ou=Admin,ou=Global Preferences,ou=[user root]”:</u> - Authenticated users - Read, Search, Compare
User root subtree (dn: [user-root])	<u>Base:</u> - Configuration Administrators group - All - administrator user account - All - Directory Administrators group - All - Server Instance Entry group - All - Authenticated users – Read, Search, Compare ** -- all attributes except “userPassword” - Self access – Write ** -- only specific attributes approved and documented - [IAO-approved groups] – [approved access] -- specific attributes approved and documented <u>“ou=People,ou=[user root]”:</u> - Self access – Write ** -- only specific attributes approved and documented - [IAO-approved groups] – [approved access] -- specific attributes approved and documented, except “cn”, “nsRoleDN”, “sn”, and “uid”

**Table E-7. Directory Database Object Access**

The following notes apply to the specifications in this table:

- The Directory Manager account is not listed in the table because that account has implicit access for all entries in the directory database.
- Except for the root DSE, anonymous access is not permitted to any part of a non-public directory. Permissions in this table represent the conversion from the vendor’s default ACIs that permit anonymous access (“anyone”) to the same access for authenticated users

(“all”). Anonymous access to the root DSE is only allowed if there is IAO-approved documentation that it is required.

- Permissions in the user root subtree that are marked with a double asterisk (\*\*) must be documented and approved by the IAO. These permissions are not required, but may be defined if necessary for the intended directory use.
- Permissions to the NetscapeRoot “ou=UserPreferences,ou=[user root]” and “cn=PublicViews,ou=4.0, ou=Admin,ou=Global Preferences,ou=[user root]” entries have been reduced from the vendor’s defaults that allow update access to authenticated users.

The specified access permissions are intended to provide a high degree of compatibility with custom, GOTS, and COTS applications. The following additional guidance is recommended:

- The SA should work to reduce the non-update (read, search, or compare ACI) permissions that are granted to all authenticated users to the following:
  - cn=SNMP,cn=config
  - Monitor subtree (cn=monitor)
  - Schema attribute (cn=schema)
  - NetscapeRoot subtree (o=NetscapeRoot).

The RHDS ACI implementation includes several flexible methods of specifying user access to directory data. Two of the bind rule criteria can be used to allow or deny access based on source. These criteria use the network addresses of the originating host of the access. The “ip” keyword is used to control access based on IP address. The “dns” keyword is used to control access based on DNS hostname. Both of these keywords allow the use of a wildcard character that allows multiple hosts to meet the criteria.

While the use of network address criteria in ACIs can be a simple and efficient way of controlling access, there are potential vulnerabilities. Address spoofing or DNS poisoning attacks can result in corrupt data being used in access control decisions. Consequently it is necessary to restrict the use of network address criteria.

- *(DS20.3120: CAT I) The IAO will ensure the use of the “ip” and “dns” keywords in ACIs is restricted as follows:*
  - *All IP addresses and host names must be under the control of a DoD Component.*
  - *All “dns” specifications must include fully qualified domain names and the DNS domains must be under the control of a DoD Component.*

The only valid exception to this requirement is for directory services managing data at the public confidentiality level. In that case the risk from invalid network address data is insignificant because the data can be made available to any user.

The “authmethod” keyword in an ACI bind rule is a method of specifying access based on the authentication method used in the bind. In the absence of other filtering criteria, the use the “none” value for “authmethod” might enable anonymous access to the directory.

- *(DS20.3130: CAT I) The IAO will ensure “none” is not specified for the “authmethod” keyword in an ACI bind rule unless other specifications will ensure authentication is performed for access to non-public data.*

The Enable Access Control configuration option provides a simple way to deactivate all access control to the directory data. When configured to “off”, all bind connections (including anonymous) have full access. While this might have value in certain public implementations where the directory is regularly rebuilt, it is not useful to DoD implementations because it would allow the directory database itself to be vulnerable to every user. Even DoD directories holding data at the public confidentiality level require control of access to the directory configuration data.

- *(DS20.3140: CAT I) The IAO will ensure the Enable Access Control (nsslapd-accesscontrol) attribute of the Configuration subtree (cn=config) is set to “on”.*

Database links are the mechanism used to implement directory server chaining. As indicated in Section E.2.3, Other Implementation Features and Details, chaining is a directory knowledge reference in which a client’s request is relayed by the original server to another server where the information resides. Because the use of a database link involves the specification of a user and associated credentials in the link definition, there are security considerations.

The nsMultiplexorBindDN attribute of a database link definition identifies the account on the server with the target data that is to be used by the server where the database link is defined. This account is referenced in the proxy access rights in the ACI that controls access to the data. To ensure appropriate access control is maintained on the target server, this account cannot have administrative-level access.

- *(DS20.3150: CAT I) The IAO will ensure the account referenced in the nsMultiplexorBindDN attribute of a database link is not the Directory Manager account or a member of the Administrators, Configuration Administrators, Directory Administrators, or Server Instance Entry group on the target directory server.*

The database links mechanism supports cascading chaining in which a first server has a link to a second server which has a link to a third server (and possibly so on). The links on servers between the first server and the final target server are called intermediate database links.

Because servers with intermediate database links create an access path to data, but do not authenticate the original client, it is necessary to restrict access (via ACI) to the target data definition on the intermediate server to the proxy user from the accessing server. The details of this configuration are described in Red Hat’s *Administrator’s Guide, Red Hat Directory Server*. However, there is a configuration control that must be set to enforce checking of the ACI. The nsCheckLocalACI attribute of a database link definition is set to “on” to enforce evaluation of the ACI.

- *(DS20.3160: CAT I) The IAO will ensure the nsCheckLocalACI attribute on all intermediate database links is set to “on”.*

A final issue with regard to database access in RHDS is related to the DSML Gateway component. If this component is installed and activated using the default configuration options, it is set to use anonymous access to the directory. To ensure an anonymous access path is not inadvertently created, a specific setting is required in the DSML Gateway configuration file.

The UseAuth configuration file operand controls whether anonymous, a user specified in the configuration file, or a user specified in each transaction is used. Anonymous access is prohibited to non-public data. Specifying a user in the configuration file must be prohibited because the password must also be specified and the configuration file is not encrypted. By specifying “true” for the UseAuth operand, it forces the use of identification data in each transaction.

- (DS20.3170: CAT I) The IAO will ensure the UseAuth operand is set to “true” in the DSML Gateway configuration file.

#### E.3.3.4. Data Change Auditing - RHDS

Section 3.3.5, Data Change Auditing, describes a generic requirement for auditing changes to directory database objects. This section identifies specific configuration settings required for RHDS to meet the generic audit requirement.

To have sufficient data to audit directory data changes, it is necessary to use three of the log files that RHDS can produce:

- The access log contains a series of records that document client connections to the directory. Collectively the records contain the client’s IP address and other session information, bind information, and operation information.
- The audit log contains records of changes to each directory database.
- The error log contains a variable level of information on errors, server startup and shutdown times, and port number usage.

The following table lists the required core server configuration (cn=config) attribute settings for the access log.

Attribute	Value
nsslapd-accesslog	[existing path and file name]
nsslapd-accesslog-level	256 [or higher]
nsslapd-accesslog-logging-enabled	on
nsslapd-accesslog-logrotationtime	1 [note below]
nsslapd-accesslog-logrotationtimeunit	day [note below]
nsslapd-accesslog-maxlogspersdir	> 5 [note below]
nsslapd-accesslog-mode	600

**Table E-8. Access Log Controls**

The following table lists the required core server configuration (cn=config) attribute settings for the audit log.

Attribute	Value
nsslapd-auditlog	<i>[existing path and file name]</i>
nsslapd-auditlog-logging-enabled	on
nsslapd-auditlog-logrotationtime	1 [note below]
nsslapd-auditlog-logrotationtimeunit	day [note below]
nsslapd-auditlog-maxlogspersdir	> 5 [note below]
nsslapd-auditlog-mode	600

**Table E-9. Audit Log Controls**

The following table lists the required core server configuration (cn=config) attribute settings for the error log.

Attribute	Value
nsslapd-errorlog	<i>[existing path and file name]</i>
nsslapd-errorlog-level	16384 [or higher]
nsslapd-errorlog-logging-enabled	On
nsslapd-errorlog-logrotationtime	1 [note below]
nsslapd-errorlog-logrotationtimeunit	day [note below]
nsslapd-errorlog-maxlogspersdir	> 5 [note below]

**Table E-10. Error Log Controls**

The following notes apply to the specifications in the log tables above:

- The ...-logrotationtime and ...-logrotationtimeunit values are intended to ensure log rotation occurs at least daily. It is acceptable to choose more frequent rotations as long as each file is captured by the backup solution used at the site.
- The ...-maxlogspersdir value is intended to ensure log file rotation does not cause a file to be deleted before backup. It is acceptable to choose any value greater than 1 as long as each file is captured by the backup solution used at the site.
- The table does not contain a value for ...-logrotationsync-enabled, but the IAO should ensure this attribute is set to "on". This enables specific log rotation times that provide better log file management and easier access when needed.

Because the impact of these settings depends on activity in the local environment, SAs are advised to monitor the log files on their directory server and make adjustments for any increase in the amount of log data.

### E.3.3.5. Group Membership and Limiting Privileges - RHDS

Section 3.3.6, Group Membership and Limiting Privileges, describes general requirements for limiting the ability to change the members of a group or change membership in groups with privileged access to the directory. This section identifies additional specific group and privilege control requirements and implementation directions that are unique for RHDS.

It is noted in Section E.2.1, Account Definitions, that some groups are automatically created when RHDS is installed and those groups have default access privileges over RHDS database objects. Membership in these groups grants access permissions to their members to change data and configuration settings that impact the entire directory.

If an unauthorized user gains membership in the Configuration Administrators, Directory Administrators, or SIE groups, that user would be able to display, add, or change RHDS objects that could impact the confidentiality, integrity, or availability of the directory.

- *(DS20.3180: CAT II) The IAO will limit the number of users and document those users assigned to the following groups: Configuration Administrators and Directory Administrators.*
- *(DS20.3190: CAT II) The IAO will ensure the SIE group includes only directory entries associated with directory server instances.*

The implementation of any RHDS directory commonly involves the creation of local groups or roles that define collections of accounts. Section E.2.1, Account Definitions, discusses the support in RHDS for static groups, dynamic groups, managed roles, and filtered roles. Some specific requirements for usage are needed to control authorization and privilege assignment based on groups or roles.

In some cases a directory administrator may define new static or dynamic groups within the NetscapeRoot and user root subtrees in order to implement administrative delegation strategies or enable group authorization schemes. While it is not possible to devise specific group structure requirements that would work in all DoD environments, a level of security is achieved by documenting and reviewing which accounts have permission to alter the group structure.

- *(DS20.3200: CAT II) The IAO will ensure locally defined users and groups that have update authority to directory entries under the “ou=Groups, ou=TopologyManagement, o=NetscapeRoot” and the “ou=groups, [user root]” entries are documented.*

Managed roles and filtered roles are the current, vendor-recommended method of implementing account collections in RHDS. Both role types can be used to allow role-based access to resources. To ensure the ability to assign roles is controlled, it is necessary to document and review the accounts that have permission to update the role definitions and the attributes that confer filtered role membership.

- *(DS20.3210: CAT II) The IAO will ensure locally defined users and groups that have update authority to the managed role and filtered role definitions are documented.*
- *(DS20.3215: CAT II) The IAO will ensure locally defined users and groups that have update authority to the nsRoleDN attribute are documented.*
- *(DS20.3220: CAT II) The IAO will ensure locally defined users and groups that have update authority to attributes that are used in filtered role definitions are documented.*
- *(DS20.3225: CAT II) The IAO will ensure non-privileged users do not have the authority to update attributes for their own account that are used in filtered role definitions.*

Note that there is one exception to the restriction on the ability of users to update their own role-related attributes. If the related role is *\*not\** used for resource access control, users may be permitted to update the attribute. For example, the implementation of an opt-in mail list could allow users to update the attribute used for mail list participation. This type of usage must be documented.

A related consideration for group and role usage is the need to enable the referential integrity mechanism in RHDS. Refer to Section E.3.3.6, Functional Configuration, for information on this subject.

### **E.3.3.6. Functional Configuration - RHDS**

Section 3.3.7, Functional Configuration, states general requirements for replication scheduling, database integrity, account lockout, and OS services. The following information identifies any additional, unique requirements and control settings needed for RHDS environments to comply with those requirements.

Replication in RHDS is described briefly in Section E.2.3.1, Replication. Replication between RHDS directory servers is configured primarily through the replication agreement defined on the supplier server. Replication scheduling is controlled through the nsDS5ReplicaUpdateSchedule attribute of each replication agreement.

When directory data is used for access control decisions and that data is replicated between servers, timely replication of changes is critical for enforcing desired access control. To comply with the daily replication requirement (DS00.3230) in section 3.3.7 for a RHDS implementation that includes replicating servers, the nsDS5ReplicaUpdateSchedule attribute in the replication agreement must be configured to allow replication for some period on all days of the week.

Note that additional considerations for replication security are included in Section E.3.3.7, Data Transmission Confidentiality and Integrity - RHDS.

Data integrity is a critical security and operational factor for all database implementations. One aspect of data integrity concerns the need for concurrent updates of interrelated information. Referential integrity is the common mechanism used to address this concern. In RHDS, static

groups and managed roles are examples of data that require the use of a referential integrity mechanism to maintain database integrity and directory security.

When an account is deleted, any static group entries that include the account (uniqueMember attribute) are not automatically updated. When a managed role is deleted, any account or nested role entries that include the managed role (nsRoleDN attribute) are not automatically updated. If an identical account or role is subsequently added, an inappropriate relationship might be created. When groups or roles are being used for access control, this could cause a serious security vulnerability.

Although the required updates could be performed manually, there is a distinct possibility that some necessary updates would be missed. The need for an automated mechanism to ensure the updates is clear. By configuring and enabling the Referential Integrity Plug-in in RHDS, this need is satisfied.

To comply with the referential integrity requirement (DS00.3240) in section 3.3.7, the RHDS Referential Integrity Plug-in must be enabled on the appropriate servers and must check the uniqueMember and nsRoleDN attributes.

**NOTE:** It is critical to consult the *Administrator's Guide, Red Hat Directory Server* document before enabling the Referential Integrity Plug-in. If the appropriate indexing is not enabled for the attributes, directory server performance can be severely degraded.

The following notes also apply to the Referential Integrity Plug-in:

- The nsRoleDN attribute is not included in the default configuration of attributes managed by the plug-in, so explicit action is needed to add it.
- The directory replication configuration dictates which servers can have the plug-in enabled. Generally it is enabled on the supplier server, not a consumer server. Consult the vendor documentation for guidance.

The general requirements related to the use of groups and roles in RHDS are discussed in Section E.3.3.5, Group Membership and Limiting Privileges - RHDS.

Compliance with the account lockout requirement (DS00.3250) in section 3.3.7 reduces the threat from password-guessing and other attacks that use invalid authentication data. If lockout is based on directory server controls (instead of or in addition to application controls), the following RHDS core server configuration (cn=config) attribute settings for the account lockout policy are required:

Attribute	Value
passwordLockout	on
passwordLockoutDuration	3600
passwordMaxFailure	3
passwordResetFailureCount	3600

**Table E-11. Password Policy Controls – Lockout**



The following notes apply to the specifications in the table above:

- These values are intended to lock an account for sixty minutes after three failed logon attempts within sixty minutes.
- These are minimum requirements that may be exceeded by local policy or when a more restrictive policy is in effect due to JTF-GNO direction or INFOCON status changes.

#### **E.3.3.7. Data Transmission Confidentiality and Integrity - RHDS**

Section 3.3.8, Data Transmission Confidentiality and Integrity, states general requirements for the protection of administrative and replication sessions over networks. The following information identifies guidance and control settings for RHDS environments to comply with those requirements.

There are multiple requirements (DS00.3280, DS00.3320) in section 3.3.8 that specify encryption for replication sessions. These requirements apply when replication occurs over wireless or non-DoD networks and when clear-text authentication data is transmitted. Although network-based encryption is a possible solution, RHDS includes the necessary support to perform application-level encryption.

Replication sessions can be encrypted through configuration steps on the supplier and consumer replication servers. The *Administrator's Guide, Red Hat Directory Server* document provides details, but essentially the following actions are necessary:

- The consumer server must be configured to support SSL using a server PKI certificate.
- The replication agreement on the supplier server must be configured (usually through the Replication Agreement Wizard) to have "SSL" in the nsDS5ReplicaTransportInfo attribute.

The architecture of RHDS affects the implementation of the requirement (DS00.3290) in section 3.3.8 for encryption of administrative sessions. This is due to the fact that there are multiple paths to accomplish administrative tasks in RHDS. One is through standard command-line programs that communicate over a network using LDAP/LDAPS. The second is through the Red Hat Console and Administration Server that is bundled with RHDS. In both cases it is possible to use network-based encryption such as what is available with a VPN or encryption features within RHDS. The following information discusses the application-based methods supported by RHDS.

To implement RHDS-based encryption of administrative sessions that use command-line programs, the following actions are necessary:

- The directory server must be configured to support SSL using a server PKI certificate.
- The client command-line programs supplied with RHDS must be invoked with the proper SSL option switches. The `ldapsearch`, `ldapmodify`, and `ldapdelete` commands offer the "-Z" option switch to specify the use of SSL. The "-ZZZ" option specifies that "Start TLS" mechanism must be successfully negotiated.

Although programs (command-line and GUI) from other sources are available with similar options, care must be taken to ensure the proper (FIPS) evaluation of the encryption implementation has been performed on the non-RHDS software.

Because the command-line programs might be used legitimately by non-privileged users accessing the directory, the use of the SSL-enabling command line options cannot be enforced through practical configuration controls on the server. Unless the server is configured to accept only SSL connections, a site policy is the only way to enforce encrypted administrative access via command-line programs.

To implement RHDS-based encryption of administrative access through the Red Hat Console and Administration Server software, the following actions are necessary:

- The directory server must be configured to support SSL using a server PKI certificate.
- The administration server must be configured to access the directory server using SSL.
- The administration server must be configured to support SSL using a server PKI certificate.
- The console connection must be configured to access the administration server using the HTTPS protocol.

There are multiple requirements (DS00.3300, DS00.3330, DS00.3340) in section 3.3.8 that specify mutual authentication for access to a directory server. The requirements apply to directory replication, administrative sessions, proxy identity use, and directory synchronization. It was noted in section 3.3.8 that network or application-level encryption can satisfy mutual authentication requirements, but there are features provided by RHDS that also enable mutual authentication.

The following actions are needed to implement RHDS-based mutual authentication for replication:

- The consumer directory server must be configured to support SSL using a server PKI certificate.
- The replication agreement on the supplier directory server must be configured to have "SSL" in the nsDS5ReplicaTransportInfo attribute.
- If the identity of the supplier directory server is based on ID/password, the replication agreement on the supplier directory server must be configured to have an account (created on the consumer server) in the nsDS5ReplicaBindDN attribute, the associated password in the nsDS5ReplicaCredentials attribute, and "SIMPLE" in the nsDS5ReplicaBindMethod attribute.
- If the identity of the supplier directory server is based on PKI certificate, the supplier directory server must be configured with a client and server PKI certificate and "SSLCLIENTAUTH" in the nsDS5ReplicaBindMethod attribute.
- If the identity of the supplier directory server is based on PKI certificate, the replication agreement on the supplier directory server must be configured to have the nsslapd-ssl-check-hostname attribute set to "on".

The following notes apply to these actions:

- Red Hat recommends that replication agreements be created using the Replication Agreement Wizard. This tool is accessed from the Directory Server Console and is documented in the *Administrator's Guide, Red Hat Directory Server* document.
- The use of certificate-based identity for the supplier directory server is the preferred implementation because it mitigates the vulnerabilities associated with passwords and reduces the threat that the Replication Manager account can be compromised.
- RHDS does not allow the supplier directory server's certificate to be a self-signed certificate.

The following actions are needed to implement RHDS-based mutual authentication for administrative, proxy-identity, and synchronization directory access:

- The directory server must be configured to support SSL using a server PKI certificate.
- The client software must be capable of establishing an SSL session with the directory server.
- The client (administrator, proxy-identity, or synchronization process) must have an account with a password on the directory server or the client must have an account with a client PKI certificate that has a certificate authority in common with the directory server.
- If the identity of the client is based on PKI certificate, the client software must be capable of supplying a PKI certificate instead of an account and password for access to the directory server.

To comply with the requirement (DS00.3310) in section 3.3.8 for a dedicated account for replication, the replication agreement on the supplier directory server must be configured to have an account in the nsDS5ReplicaBindDN attribute that is not used for any other purpose on the consumer directory server. However, if the replication identity for the supplier directory server is based on PKI certificate, the nsDS5ReplicaBindDN attribute must be empty or not specified.

Directory service configurations that use certificate-based authentication have to comply with the requirement (DS00.3350) in section 3.3.8 to incorporate CRL checking. Both the directory server and the Administration Server may be configured to use SSL. Each server that uses SSL must be configured for CRL checking. The following actions are needed:

- The directory server must be updated regularly with the current CRL(s) for each DoD CA defined to the server.
- The administration server must be updated regularly with the current CRL(s) for each DoD CA defined to the server.

The following notes apply to these actions:

- The "Manage Certificates" task in the RHDS console and the administration server console can be used to import CRLs.
- The frequency at which CRLs are updated must be determined by the directory owner, dependent on the security requirements for the data. In general, the CRL for a directory service that is used to control access to resources at the sensitive confidentiality level and above should be updated at least daily, and must be updated at least weekly.
- Refer to the current JTF-GNO documents for guidance related to CRL download issues and efficient PKI certificate validation approaches.

It is possible to configure RHDS to implement the requirement (DS00.3370) in section 3.3.8 to limit inactive client connections. RHDS also supports different limits for individual accounts. The following configuration settings are used to implement the limits:

- The `nsslapd-idletimeout` core server configuration (`cn=config`) attribute is used to specify the default level timeout. A value of “900” must be used to indicate that LDAP client connections idle for 15 minutes will be closed.
- The `nsIdleTimeout` attribute can be set at the account (bind DN) level to override the default limit for documented circumstances that require it.

### E.3.4. Continuity

Section 3.6, Continuity, identifies general requirements to support backup and restore processes for directory services. The following information provides guidance on meeting the general requirements and an additional requirement that is unique to RHDS for these processes.

To comply with the directory database backup requirement (DS00.0160) in section 3.6, it is necessary to perform the following tasks on an RHDS server:

- Execute the `ns-slapd` utility (commonly through the `db2bak` script) or the directory console backup task to create a coherent backup copy of the directory database files.
- Execute the `ns-slapd` utility (commonly through the `saveconfig` script) to create an LDIF version of the directory configuration data in the `NetscapeRoot` subtree.
- If the directory server or administration server is configured for SSL, use an OS or third party backup utility to capture the PKI certificate and key databases.

The frequency at which the backup tasks must run varies according to the use and volatility of the data. The database files should generally be captured on a daily basis. The configuration (`NetscapeRoot`) data and PKI databases should be captured on a weekly basis.

The unlimited database access available to the Directory Manager account may be needed in some recovery scenarios. When recovery is performed after a physical loss or when the normal administrative personnel are not available, physical access to the password of the Directory Manager account may be the only means by which an orderly recovery can be completed.

- *(DS20.6100: CAT II) The IAO will ensure the password of the Directory Manager account for each RHDS directory server is stored in a locked, fire-rated container or is subject to other appropriate protections from loss.*

To recover an RHDS directory service that includes multiple servers with replication configurations, it is necessary to know the placement of each server within the directory implementation architecture. This placement refers to the role of each server as supplier or consumer in the context of replication.

To comply with the directory architecture details requirement (DS00.6120) in section 3.6, the following information is needed:

- Text or preferably a graphic representation of the logical directory architecture that lists the servers (directory and administration) by host and directory server name, and indicates the replication (supplier \ consumer) flow

- Indications of which directory servers are configured to support SSL
- The port numbers for the directory and administration servers
- The OS accounts (users and groups) under which each directory server and administration server execute
- The bind DN of the Directory Manager account for each server
- The suffixes (naming contexts) covered by the directory.

When used with the proper backup media, this collection of information should be sufficient to restore the directory service to operating status.

This page is intentionally left blank.